AI-Human Collaboration Levels Up › Mathematicians formalize proofs with AI
P.5

How to Fight an Internet Blackout › A bold scheme hides data in satellite streams
P.38

Computing Over the Air › Building wireless networks that exploit radio interference
P.44

FOR THE TECHNOLOGY INSIDER
MAY 2026

IEEE Spectrum

WANDERCRAFT
E 26 St
Herman Melville Square
ONE WAY

Exoskeleton Test Pilot Robert Woo Pushes the Limits

IEEE

There's more to LTC than just
where you receive care.

Expertise when your loved ones need help.

1-800-493-IEEE (4333)
To learn more, visit IEEEinsurance.com/IEEELTC

IEEE
LONG-TERM CARE

Program Administered by AMBA Administrators, Inc.
In CA d/b/a Association Member Benefits & Insurance Agency
CA LIC #0I96562 • AR License #100114462

107214 (5/26) Copyright 2026 AMBA. All rights reserved

VOLUME 63 / ISSUE 5 MAY 2026

Contents

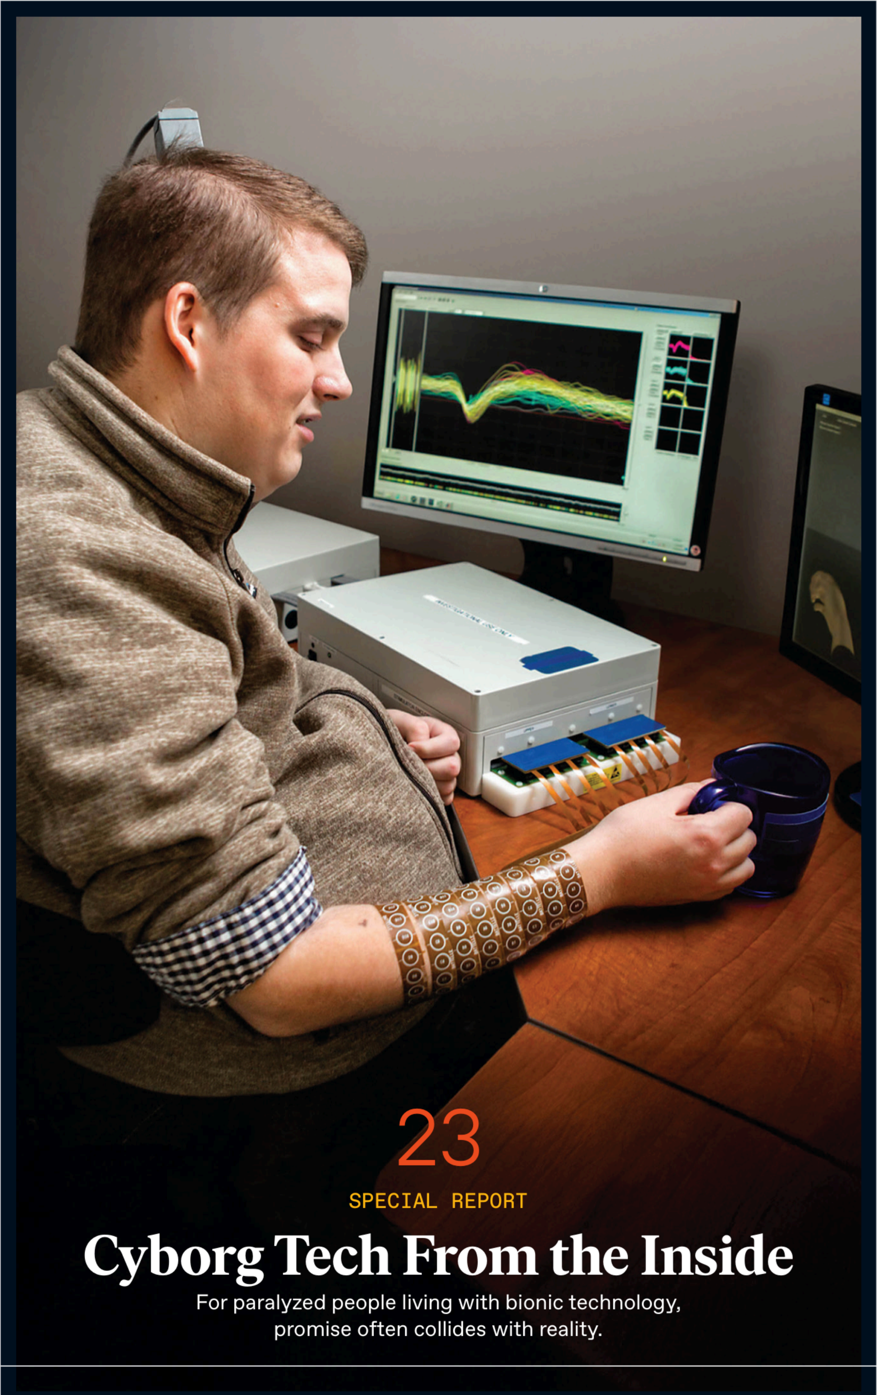

BATTELLE

23
SPECIAL REPORT
Cyborg Tech From the Inside
For paralyzed people living with bionic technology, promise often collides with reality.

The Man in the Machine 24
One user spent 15 years pushing exoskeletons to improve.
By Eliza Strickland

Life With an Experimental Brain Implant 32
Early users reveal what brain implants demand and deliver.
By Edd Gent

The Stealth Signals Bypassing Iran's Internet Blackout 38
Despite the information shutdown, data hidden in satellite streams is still getting through.
By Evan Alireza Firoozi

Making Radio Waves Compute in Midair 44
This tech turns interference into a chipless calculator.
By Ana I. Pérez-Neira & Alphan Sahin

EDITOR'S NOTE 2
The bionic age relies on the ultimate beta testers.

NEWS 5
AI Does Math
Data Centers at Sea
Self-Driving on Mars

THE DATA 14
GPU renters are playing a silicon lottery.

HANDS ON 16
Putting teletext back on the airwaves.

CAREERS 20
This teacher sees neurodiversity as a superpower.

5 QUESTIONS 22
Jon Wagner designs eVTOL motors.

PAST FORWARD 52
Hershey's Electric Railroad in Cuba

ON THE COVER:
Photo by Bryan Anselm/Redux

BY ELIZA STRICKLAND

Beyond the Demo

Exoskeletons and brain-computer interfaces may seem miraculous, but their true value is forged in the real world

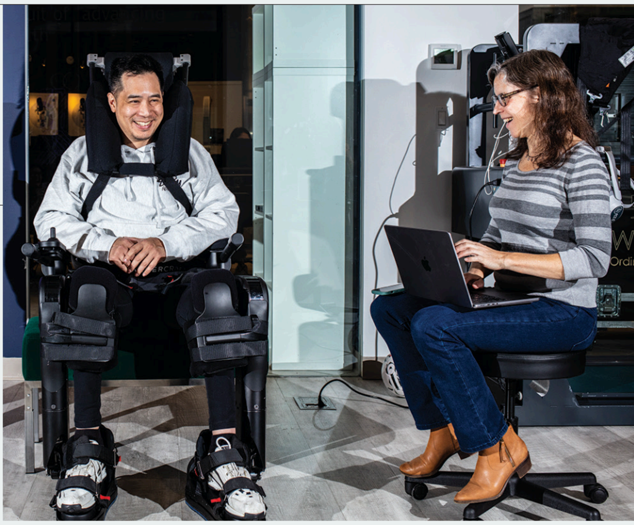

Robert Woo sits in a self-balancing exoskeleton during an interview with *IEEE Spectrum* Features Editor Eliza Strickland.

I first met Robert Woo in 2011, during his third time walking in a powered exoskeleton. The architect had been paralyzed in a construction accident four years earlier, but he was determined to get back on his feet. Watching him clunk across a rehab room in an exoskeleton prototype, the technology felt astonishing. I had the same reaction when reporting on early brain-computer interfaces (BCIs), which enabled paralyzed people to move robotic arms or communicate by thought alone. Both types of bionic technology seemed to verge on magic.

But that initial sense of awe, as I've learned over many years of reporting on these technologies, is only a starting point. What matters is not what these systems can do in a carefully staged demo but how they perform in the real world. Do they work reliably? Can people with disabilities use them for their intended purposes? And what does it actually cost—in time, effort, and trade-offs—to do so? The question isn't whether the technology looks impressive the first time but whether it holds up on the hundredth.

The special report in this issue, "Cyborg Tech From the Inside" [p. 23], takes that perspective seriously. In my feature article on Woo [p. 24], an exoskeleton super-user who has spent 15 years testing these systems, the story of the technology is inseparable from the story of its use. Woo's relentless feedback has driven steady, incremental improvements. In Edd Gent's reporting on the pioneers testing the earliest BCIs [p. 32], the experience of these extraordinary technologies likewise resolves into something more complex. As one trial participant notes, these early adopters are like the first astronauts, who barely reached space before coming back down to Earth. Together, these stories reframe these individuals not as passive medical patients but as the ultimate beta testers and co-engineers of the bionic age.

These users aren't just patients—they're the beta testers and co-engineers of the bionic age.

I saw the gap between demonstration and daily use firsthand when I interviewed Woo in a Manhattan showroom recently, where he was testing a new self-balancing exoskeleton from Wandercraft. The device is a striking advance that kept him upright without crutches, but it also revealed the friction of the real world. As Woo tried to walk out the door, barely an inch of slope on the Park Avenue sidewalk was enough to trigger the machine's safety sensors and halt his progress. It was a stark reminder of how far these systems must evolve before they fit seamlessly into everyday life.

For the people who use them, that seamless integration is the ultimate goal. Getting there will depend not just on technical breakthroughs but on how well these systems hold up outside controlled environments, over time, and under real conditions. Looking from the inside doesn't make these technologies any less remarkable, but it does change how we judge them—not by what they can do once for a photo but by what they can sustain over a lifetime. That's the standard their users have been applying all along.

Our commitment to evaluating technology from the user's perspective extends beyond this special report. To provide a necessary corrective to the "techno-solutionism" that often dominates coverage of assistive devices, *IEEE Spectrum* created the Taenzer Fellowship for Disability-Engaged Journalism, under which six writers with disabilities are contributing articles about the devices they rely on daily. As Special Projects Director Stephen Cass notes, these journalists "aren't afraid to ask clear-eyed questions about the tech and are deeply aware of how it impacts humans." You can read the fellows' work at spectrum.ieee.org/tag/taenzer-fellowship. ■

LEFT: NICOLE MILLMAN

Photo by Bryan Anselm/Redux

ROBOT OF THE MONTH

FURHAT

Furhat is a conversational robot head with a back-projected face designed to interact with humans in social settings. With highly expressive facial animations and a vast library of voices, it is used in academic research, education, and commercial innovation around the world.

CREATOR
Furhat Robotics

YEAR
2018

COUNTRY
Sweden

CATEGORY
Education, Research

FEATURES
AI, full autonomy, speech, vision

PHOTO: FURHAT ROBOTICS

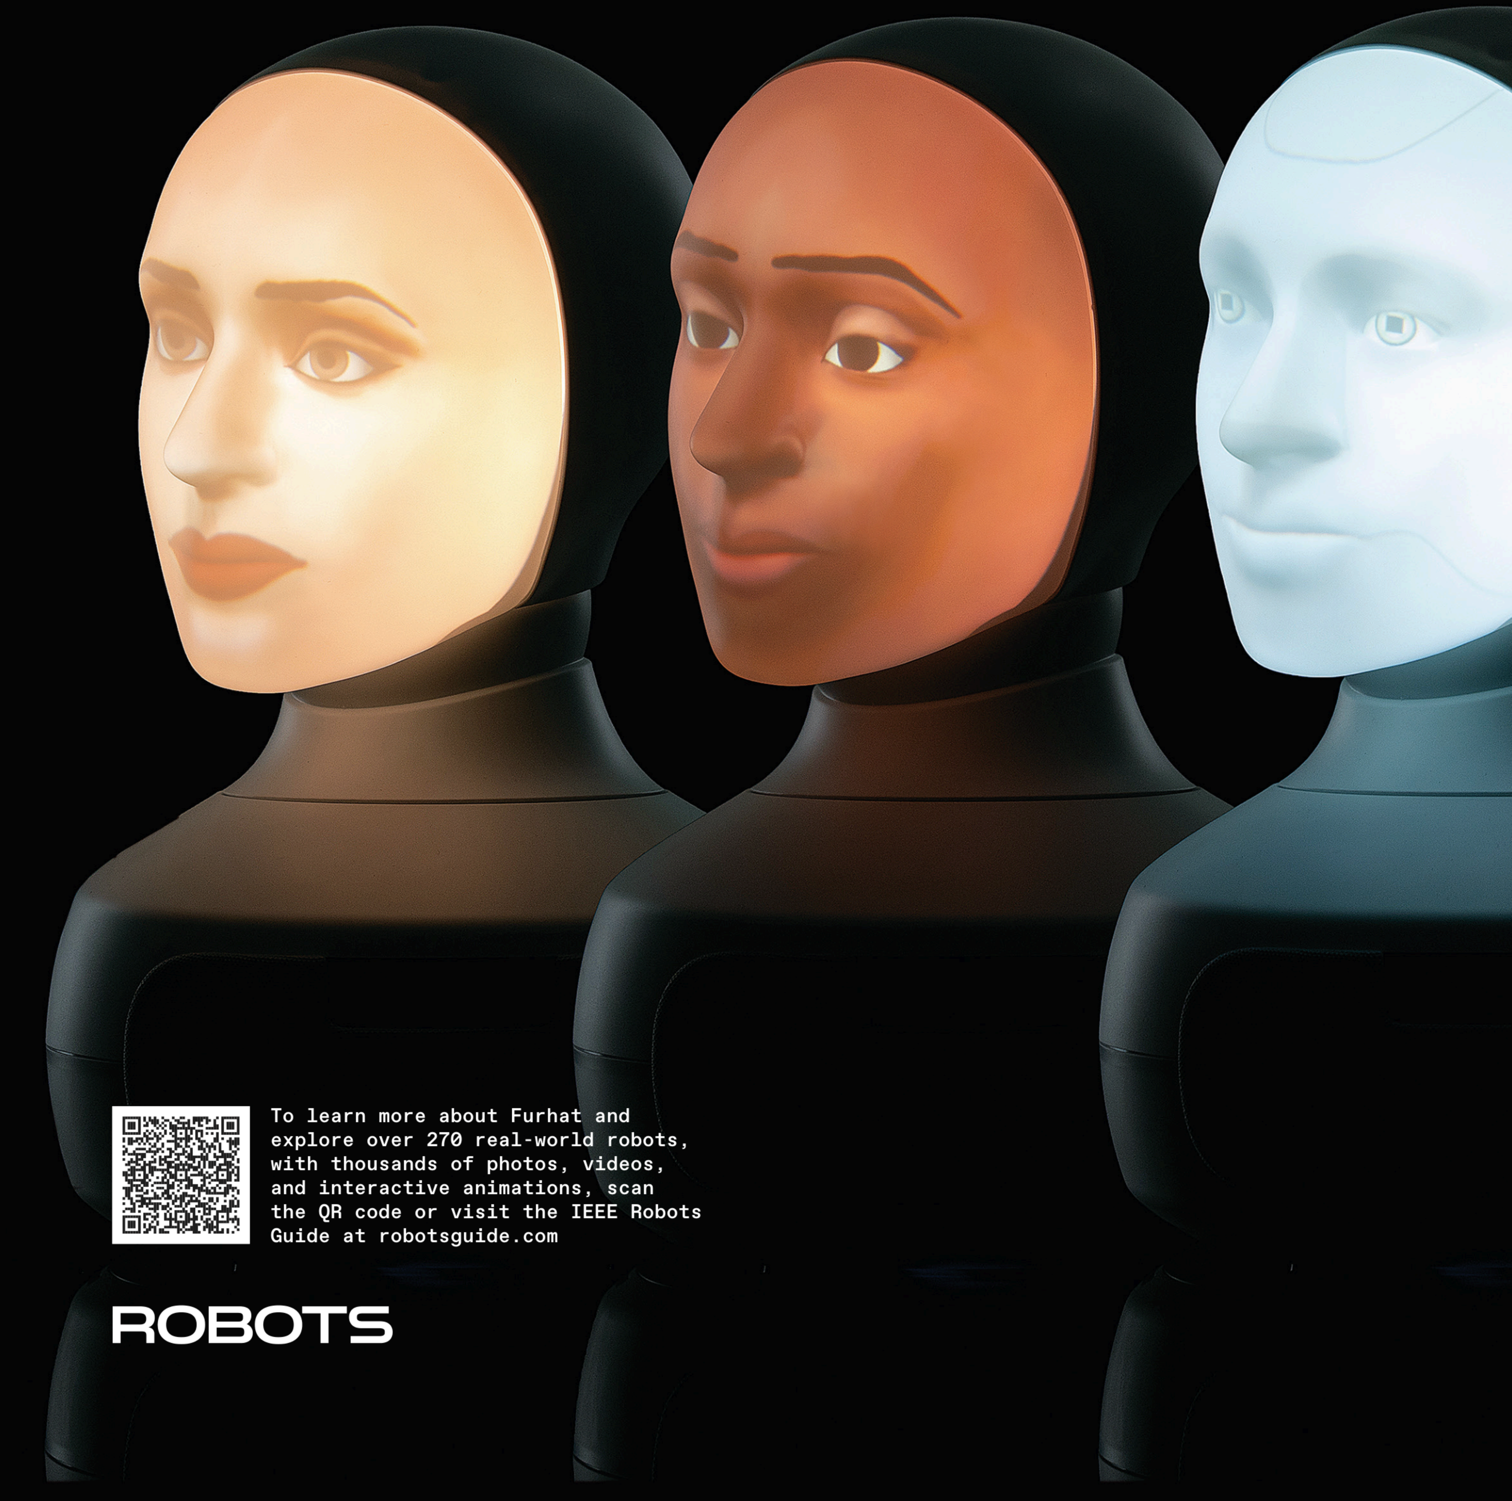

To learn more about Furhat and explore over 270 real-world robots, with thousands of photos, videos, and interactive animations, scan the QR code or visit the IEEE Robots Guide at robotsguide.com

ROBOTS

EVAN ALIREZA FIROOZI

Firoozi is the executive director of NetFreedom Pioneers, an expert on information accessibility, and an advocate for internet freedom. In this issue, he writes about NetFreedom's Toosheh system, which sends news and other data via satellite streaming [p. 38]. Such untraceable transmissions have been a vital source of information during Iran's internet blackout. Born in Iran, Firoozi grew up understanding the consequences of government censorship. "From Iran to beyond, access to information is not a luxury," he says. "It's a lifeline."

EDD GENT

Based in Bengaluru, India, Gent writes about emerging technologies across computing, engineering, and bioscience. While reporting his story about paralyzed users of experimental brain implants [p. 32], he was struck by what mattered most to them. "I had assumed the most life-changing benefits of this technology would be around autonomy and independence," he says. But users said the biggest impacts were "less practical and more emotional," emphasizing things like stroking a pet or playing a video game.

ANA I. PÉREZ-NEIRA

Pérez-Neira is director of Centre Tecnològic de Telecomunicacions de Catalunya and a professor at the Universitat Politècnica de Catalunya, in Barcelona. Coauthor Alphan Sahin is an assistant professor of electrical engineering at the University of South Carolina in Columbia, S.C. In this issue [p. 44], they write about over-the-air computation—"an idea that sits on the cusp of hype and breakthrough," Sahin says. Pérez-Neira adds, "Our work asks a simple question: What if signal channels aren't just for communication but for computation too?"

BENJAMIN SKUSE

Skuse, a freelance writer with a Ph.D. in applied mathematics, writes about science and technology. In this issue, he explains how AI models contribute to the formal verification of mathematical theorems [p. 5]. Working on the story opened his eyes to how quickly reasoning AI agents are improving, he says: "These rapid advances make me wonder how far we might be from AI independently solving some of the biggest problems in math, and what that might mean for human mathematicians."

IEEE Spectrum

EDITOR IN CHIEF Harry Goldstein, h.goldstein@ieee.org
EXECUTIVE EDITOR Jean Kumagai, j.kumagai@ieee.org
MANAGING EDITOR Elizabeth A. Bretz, e.bretz@ieee.org
CREATIVE DIRECTOR Mark Montgomery, m.montgomery@ieee.org
DIRECTOR OF DIGITAL INNOVATION Erico Guizzo, e.guizzo@ieee.org
EDITORIAL DIRECTOR, CONTENT DEVELOPMENT
Glenn Zorpette, g.zorpette@ieee.org

FEATURES EDITOR Eliza Strickland, e.strickland@ieee.org

SPECIAL PROJECTS DIRECTOR Stephen Cass, cass.s@ieee.org

SENIOR EDITORS
Evan Ackerman, evan.ackerman@ieee.org
Lucas Laursen, l.laursen@ieee.org
Samuel K. Moore, s.k.moore@ieee.org

NEWS MANAGER Michael Koziol, m.koziol@ieee.org

SENIOR ASSOCIATE EDITORS
Margo Anderson, m.k.anderson@ieee.org
Emily Waltz, e.waltz@ieee.org

ASSOCIATE EDITORS
Dina Genkina, d.genkina@ieee.org
Willie D. Jones, w.jones@ieee.org

ASSISTANT EDITOR Gwendolyn Rak, g.rak@ieee.org

PRODUCT AND AUDIENCE DEVELOPMENT DIRECTOR
Tim Warder, t.warder@ieee.org
PRODUCT MANAGER Priyanka Shitole, p.shitole@ieee.org
AUDIENCE DEVELOPMENT MANAGER
Laura Bridgeman, l.bridgeman@ieee.org
COMMUNITY MANAGER
Kohava Mendelsohn, k.mendelsohn@ieee.org
DIGITAL PRODUCT DESIGNER Erik Vrielink, e.vrielink@ieee.org

ART & PRODUCTION
ART DIRECTOR Brandon Palacio, b.palacio@ieee.org
PHOTOGRAPHY DIRECTOR Randi Klett, randi.klett@ieee.org
ASSISTANT PHOTO EDITOR Nicole Millman, n.millman@ieee.org
PRINT PRODUCTION SPECIALIST
Sylvana Meneses, s.meneses@ieee.org
DIGITAL PRODUCTION MANAGER
Michael Spector, m.spector@ieee.org

SENIOR COPY EDITOR Joseph N. Levine, j.levine@ieee.org
COPY EDITOR Michele Kogon, m.kogon@ieee.org
RESEARCH COPY EDITOR Cheryl Della Pietra, c.dellapietra@ieee.org

ADMINISTRATOR Jose Zambrano, jose.zambrano@ieee.org

CONTRIBUTING EDITORS Rina Diane Caballar, Robert N. Charette, Charles Q. Choi, Tom Clynes, Peter Fairley, Edd Gent, W. Wayt Gibbs, Mark Harris, Allison Marsh, Julianne Pepitone, Ned Potter, Matthew S. Smith, Lawrence Ulrich

THE INSTITUTE
EDITOR IN CHIEF Kathy Pretz, k.pretz@ieee.org
ASSOCIATE EDITOR Joanna Goodrich, j.goodrich@ieee.org

SUBSCRIPTION SERVICES
IEEE Members member-services@ieee.org
Non-members +1 866 363 2304

DIRECTOR, PERIODICALS PRODUCTION SERVICES Peter Tuohy
ADVERTISING PRODUCTION MANAGER
Felicia Spagnoli, f.spagnoli@ieee.org

ADVERTISING PRODUCTION +1 732 562 6334

EDITORIAL ADVISORY BOARD, *IEEE SPECTRUM*
Harry Goldstein, *Chair*; Muhammad Alaa, Karanveer Anand, Ella M. Atkins, Sangyeun Cho, Hugh Durrant-Whyte, Benjamin Gross, Carlos Gutierrez, Lawrence O. Hall, Benjamin Kroposki, Somdeb Majumdar, Khaled Mokhtar, Ramune Nagisetty, Paul Nielsen, Ketra Schmidt, Gaurav Sharma, Fan Shi, Mini S. Thomas, Andrej Zdravkovic, Lidong Zhou

EDITORIAL ADVISORY BOARD, *THE INSTITUTE*
Kathy Pretz, *Chair*; Stamatis Dragoumanos, Madeleine Glick, Harry Goldstein, Lawrence O. Hall, Francesca Iacopi, Sai Prashanth Mallellu, Shashi Raj Pandey, John Purvis, Chenyang Xu

EDITORIAL CORRESPONDENCE
IEEE Spectrum, 3 Park Ave., 17th Floor,
New York, NY 10016-5997 Tel: +1 212 419 7555

ADVERTISING INQUIRIES WILEY
Vanessa Winde +49 620 160 6721, vwinde@wiley.com
Chris Darch +1 206 867 4165, cdarch@wiley.com

REPRINT SALES +1 212 221 9595, ext. 319

REPRINT PERMISSION/ LIBRARIES Articles may be photocopied for private use of patrons. A per-copy fee must be paid to the Copyright Clearance Center, 29 Congress St., Salem, MA 01970. For other copying or republication, contact Managing Editor, *IEEE Spectrum.*

COPYRIGHTS AND TRADEMARKS *IEEE Spectrum* is a registered trademark owned by The Institute of Electrical and Electronics Engineers Inc. Responsibility for the substance of articles rests upon the authors, not IEEE, its organizational units, or its members. Articles do not represent official positions of IEEE. Readers may post comments online; comments may be excerpted for publication. IEEE reserves the right to reject any advertising.

IEEE

IEEE BOARD OF DIRECTORS
PRESIDENT&CEO Mary Ellen Randall, president@ieee.org
+1 732 562 3928 Fax: +1 732 981 9515
PRESIDENT-ELECT Jill I. Gostin
PAST PRESIDENT Kathleen A. Kramer
TREASURER Gerardo Barbosa
SECRETARY Jenifer Castillo Rodriguez

VICE PRESIDENTS
Timothy P. Kurzweg, *Educational Activities;* Andrew D. Lowery, *Member & Geographic Activities;* W. Clem Karl, *Publication Services & Products;* F.D. Tan, *Technical Activities;* Gary R. Hoffman, *President, Standards Association;* Barry C. Tilton, *President, IEEE-USA*

DIVISION DIRECTORS
Fernando J. Guarin (I); H. Alan Mantooth (II); Sarah Kate Wilson (III); Charles M. Jackson (IV); Leila De Floriani (V); Russell D. Meier (VI); Noel M. Schulz (VII); Cecilia Metra (VIII); Cristina Olaverri-Monreal (IX); Ravinder S. Dahiya (X)

REGION DIRECTORS
Charles P. Rubenstein (1); Felicia Reinhart (2); Sonya L. Dillard (3); Constance A. Kelly (4); Christopher B. Sanderson (5); Joseph C. Wei (6); Wahab Almuhtadi (7); Dusanka Boskovic (8); Jose-Ignacio Castillo-Velazquez (9); Takako Hashimoto (10)

IEEE STAFF
EXECUTIVE DIRECTOR & COO Sophia A. Muirhead
+1 732 562 5400, s.muirhead@ieee.org
CHIEF FINANCIAL OFFICER Kelly Armstrong
+1 732 562 6011, k.armstrong@ieee.org
MANAGING DIRECTOR,TECHNICAL ACTIVITIES Mojdeh Bahar
+1 732 562 3850, m.bahar@ieee.org
GENERAL COUNSEL & CHIEF COMPLIANCE OFFICER
Ahsaki Benion +1 212 705 8975, a.benion@ieee.org
MANAGING DIRECTOR, PUBLICATIONS Steven Heffner
+1 212 705 8958, s.heffner@ieee.org
CHIEF GOVERNANCE OFFICER Donna Hourican
+1 732 562 6330, d.hourican@ieee.org
MANAGING DIRECTOR, CONFERENCES, EVENTS AND EXPERIENCES Marie Hunter
+ 1 732 465 5889, m.m.hunter@ieee.org
MANAGING DIRECTOR, MEMBER & GEOGRAPHIC ACTIVITIES
Cecelia Jankowski +1 732 562 5504, c.jankowski@ieee.org
MANAGING DIRECTOR, EDUCATIONAL ACTIVITIES
Jamie Moesch +1 732 562 5514, j.moesch@ieee.org
CHIEF MARKETING OFFICER Jayne O'Brien
+1 732 981 3433, jayne.obrien@ieee.org
MANAGING DIRECTOR, IEEE STANDARDS Alpesh Shah
+1 732 465 6467, alpesh.shah@ieee.org
CHIEF INFORMATION AND DIGITAL OFFICER Jeff Strohschein
+1 732 562 5332, j.strohschein@ieee.org
CHIEF OF STAFF Allison Taylor
+1 732 562 6068, allison.taylor@ieee.org
CHIEF HUMAN RESOURCES OFFICER Cheri N. Collins Wideman
+1 732 562 5335, c.wideman@ieee.org
MANAGING DIRECTOR, IEEE-USA Russell T. Harrison
+1 202 530 8326, r.t.harrison@ieee.org

IEEE PUBLICATION SERVICES & PRODUCTS BOARD
W. Clem Karl, *Chair*; John T. Barr, Nazanin Bassiri-Gharb, Sergio Benedetto, Eddie Custovic, Stephen Dukes, Alessandro Foi, Lawrence O. Hall, Charles Hansen, Warunika Hippola, James Irvine, Michael Jensen, Athanasios Kakarountas, James Keller, Imran Mehdi, Khaled Mokhtar, Paolo Montuschi, Zorana Popovic, Mehrdad Saif, Gianluca Setti, Gaurav Sharma, Sarah Spurgeon, Jaideep Vaidya, Maria Elena Valcher, Patrick Wheeler, Steve Yurkovich

IEEE OPERATIONS CENTER
445 Hoes Lane, Box 1331, Piscataway, NJ 08854-1331 U.S.A.
Tel: +1 732 981 0060 Fax: +1 732 981 1721

IEEE SPECTRUM (ISSN 0018-9235) is published monthly by The Institute of Electrical and Electronics Engineers, Inc. All rights reserved. © 2026 by The Institute of Electrical and Electronics Engineers, Inc., 3 Park Avenue, New York, NY 10016-5997, U.S.A. Volume No. 63, Issue No. 5. The editorial content of *IEEE Spectrum* magazine does not represent official positions of the IEEE or its organizational units. Canadian Post International Publications Mail (Canadian Distribution) Sales Agreement No. 40013087. Return undeliverable Canadian addresses to: Circulation Department, *IEEE Spectrum*, Box 1051, Fort Erie, ON L2A 6C7. Cable address: ITRIPLEE. Fax: +1 212 419 7570. INTERNET: spectrum@ieee.org. ANNUAL SUBSCRIPTIONS: United States and Canada: $74.95. Elsewhere: $99.95. POSTMASTER: Please send address changes to *IEEE Spectrum*, ℅ Coding Department, IEEE Service Center, 445 Hoes Lane, Box 1331, Piscataway, NJ 08855. Periodicals postage paid at New York, NY, and additional mailing offices. Canadian GST #125634188. Printed at 3401 N. Heartland Dr., Liberty, MO 64068, U.S.A. Printed at Spenta Multimedia Pvt. Ltd., Plot 15, 16 & 21/1, Village Chikhloli, Morivali, MIDC, Ambernath (West), Dist. Thane. *IEEE Spectrum* circulation is audited by BPA Worldwide. *IEEE Spectrum* is a member of the Association of Business Information & Media Companies, the Association of Magazine Media, and Association Media & Publishing. IEEE prohibits discrimination, harassment, and bullying. For more information, visit https://www.ieee.org/about/corporate/governance/p9-26.html.

CONTRIBUTORS

● EVAN ALIREZA FIROOZI

Firoozi is the executive director of NetFreedom Pioneers, an expert on information accessibility, and an advocate for internet freedom. In this issue, he writes about NetFreedom's Toosheh system, which sends news and other data via satellite streaming [p. 38]. Such untraceable transmissions have been a vital source of information during Iran's internet blackout. Born in Iran, Firoozi grew up understanding the consequences of government censorship. "From Iran to beyond, access to information is not a luxury," he says. "It's a lifeline."

● EDD GENT

Based in Bengaluru, India, Gent writes about emerging technologies across computing, engineering, and bioscience. While reporting his story about paralyzed users of experimental brain implants [p. 32], he was struck by what mattered most to them. "I had assumed the most life-changing benefits of this technology would be around autonomy and independence," he says. But users said the biggest impacts were "less practical and more emotional," emphasizing things like stroking a pet or playing a video game.

● ANA I. PÉREZ-NEIRA

Pérez-Neira is director of Centre Tecnològic de Telecomunicacions de Catalunya and a professor at the Universitat Politècnica de Catalunya, in Barcelona. Coauthor Alphan Sahin is an assistant professor of electrical engineering at the University of South Carolina in Columbia, S.C. In this issue [p. 44], they write about over-the-air computation—"an idea that sits on the cusp of hype and breakthrough," Sahin says. Pérez-Neira adds, "Our work asks a simple question: What if signal channels aren't just for communication but for computation too?"

● BENJAMIN SKUSE

Skuse, a freelance writer with a Ph.D. in applied mathematics, writes about science and technology. In this issue, he explains how AI models contribute to the formal verification of mathematical theorems [p. 5]. Working on the story opened his eyes to how quickly reasoning AI agents are improving, he says: "These rapid advances make me wonder how far we might be from AI independently solving some of the biggest problems in math, and what that might mean for human mathematicians."

IEEE Spectrum

EDITOR IN CHIEF Harry Goldstein, h.goldstein@ieee.org
EXECUTIVE EDITOR Jean Kumagai, j.kumagai@ieee.org
MANAGING EDITOR Elizabeth A. Bretz, e.bretz@ieee.org
CREATIVE DIRECTOR Mark Montgomery, m.montgomery@ieee.org
DIRECTOR OF DIGITAL INNOVATION Erico Guizzo, e.guizzo@ieee.org
EDITORIAL DIRECTOR, CONTENT DEVELOPMENT
Glenn Zorpette, g.zorpette@ieee.org
FEATURES EDITOR Eliza Strickland, e.strickland@ieee.org
SPECIAL PROJECTS DIRECTOR Stephen Cass, cass.s@ieee.org
SENIOR EDITORS
Evan Ackerman, evan.ackerman@ieee.org
Lucas Laursen, l.laursen@ieee.org
Samuel K. Moore, s.k.moore@ieee.org
NEWS MANAGER Michael Koziol, m.koziol@ieee.org
SENIOR ASSOCIATE EDITORS
Margo Anderson, m.k.anderson@ieee.org
Emily Waltz, e.waltz@ieee.org
ASSOCIATE EDITORS
Dina Genkina, d.genkina@ieee.org
Willie D. Jones, w.jones@ieee.org
ASSISTANT EDITOR Gwendolyn Rak, g.rak@ieee.org
PRODUCT AND AUDIENCE DEVELOPMENT DIRECTOR
Tim Warder, t.warder@ieee.org
PRODUCT MANAGER Priyanka Shitole, p.shitole@ieee.org
AUDIENCE DEVELOPMENT MANAGER
Laura Bridgeman, l.bridgeman@ieee.org
COMMUNITY MANAGER
Kohava Mendelsohn, k.mendelsohn@ieee.org
DIGITAL PRODUCT DESIGNER Erik Vrielink, e.vrielink@ieee.org
ART & PRODUCTION
ART DIRECTOR Brandon Palacio, b.palacio@ieee.org
PHOTOGRAPHY DIRECTOR Randi Klett, randi.klett@ieee.org
ASSISTANT PHOTO EDITOR Nicole Millman, n.millman@ieee.org
PRINT PRODUCTION SPECIALIST
Sylvana Meneses, s.meneses@ieee.org
DIGITAL PRODUCTION MANAGER
Michael Spector, m.spector@ieee.org
SENIOR COPY EDITOR Joseph N. Levine, j.levine@ieee.org
COPY EDITOR Michele Kogon, m.kogon@ieee.org
RESEARCH COPY EDITOR Cheryl Della Pietra, c.dellapietra@ieee.org
ADMINISTRATOR Jose Zambrano, jose.zambrano@ieee.org
CONTRIBUTING EDITORS Rina Diane Caballar, Robert N. Charette, Charles Q. Choi, Tom Clynes, Peter Fairley, Edd Gent, W. Wayt Gibbs, Mark Harris, Allison Marsh, Julianne Pepitone, Ned Potter, Matthew S. Smith, Lawrence Ulrich
THE INSTITUTE
EDITOR IN CHIEF Kathy Pretz, k.pretz@ieee.org
ASSOCIATE EDITOR Joanna Goodrich, j.goodrich@ieee.org
SUBSCRIPTION SERVICES
IEEE Members member-services@ieee.org
Non-members +1 866 363 2304
DIRECTOR, PERIODICALS PRODUCTION SERVICES Peter Tuohy
ADVERTISING PRODUCTION MANAGER
Felicia Spagnoli, f.spagnoli@ieee.org
ADVERTISING PRODUCTION +1 732 562 6334
EDITORIAL ADVISORY BOARD, *IEEE SPECTRUM*
Harry Goldstein, *Chair*; Muhammad Alaa, Karanveer Anand, Ella M. Atkins, Sangyeun Cho, Hugh Durrant-Whyte, Benjamin Gross, Carlos Gutierrez, Lawrence O. Hall, Benjamin Kroposki, Somdeb Majumdar, Khaled Mokhtar, Ramune Nagisetty, Paul Nielsen, Ketra Schmidt, Gaurav Sharma, Fan Shi, Mini S. Thomas, Andrej Zdravkovic, Lidong Zhou
EDITORIAL ADVISORY BOARD, *THE INSTITUTE*
Kathy Pretz, *Chair*; Stamatis Dragoumanos, Madeleine Glick, Harry Goldstein, Lawrence O. Hall, Francesca Iacopi, Sai Prashanth Mallellu, Shashi Raj Pandey, John Purvis, Chenyang Xu
EDITORIAL CORRESPONDENCE
IEEE Spectrum, 3 Park Ave., 17th Floor,
New York, NY 10016-5997 Tel: +1 212 419 7555
ADVERTISING INQUIRIES WILEY
Vanessa Winde +49 620 160 6721, vwinde@wiley.com
Chris Darch +1 206 867 4165, cdarch@wiley.com
REPRINT SALES +1 212 221 9595, ext. 319
REPRINT PERMISSION/ LIBRARIES Articles may be photocopied for private use of patrons. A per-copy fee must be paid to the Copyright Clearance Center, 29 Congress St., Salem, MA 01970. For other copying or republication, contact Managing Editor, *IEEE Spectrum.*
COPYRIGHTS AND TRADEMARKS *IEEE Spectrum* is a registered trademark owned by The Institute of Electrical and Electronics Engineers Inc. Responsibility for the substance of articles rests upon the authors, not IEEE, its organizational units, or its members. Articles do not represent official positions of IEEE. Readers may post comments online; comments may be excerpted for publication. IEEE reserves the right to reject any advertising.

IEEE

IEEE BOARD OF DIRECTORS
PRESIDENT&CEO Mary Ellen Randall, president@ieee.org
+1 732 562 3928 Fax: +1 732 981 9515
PRESIDENT-ELECT Jill I. Gostin
PAST PRESIDENT Kathleen A. Kramer
TREASURER Gerardo Barbosa
SECRETARY Jenifer Castillo Rodriguez
VICE PRESIDENTS
Timothy P. Kurzweg, *Educational Activities;* Andrew D. Lowery, *Member & Geographic Activities;* W. Clem Karl, *Publication Services & Products;* F.D. Tan, *Technical Activities;* Gary R. Hoffman, *President, Standards Association;* Barry C. Tilton, *President, IEEE-USA*
DIVISION DIRECTORS
Fernando J. Guarin (I); H. Alan Mantooth (II); Sarah Kate Wilson (III); Charles M. Jackson (IV); Leila De Floriani (V); Russell D. Meier (VI); Noel M. Schulz (VII); Cecilia Metra (VIII); Cristina Olaverri-Monreal (IX); Ravinder S. Dahiya (X)
REGION DIRECTORS
Charles P. Rubenstein (1); Felicia Reinhart (2); Sonya L. Dillard (3); Constance A. Kelly (4); Christopher B. Sanderson (5); Joseph C. Wei (6); Wahab Almuhtadi (7); Dusanka Boskovic (8); Jose-Ignacio Castillo-Velazquez (9); Takako Hashimoto (10)
IEEE STAFF
EXECUTIVE DIRECTOR & COO Sophia A. Muirhead
+1 732 562 5400, s.muirhead@ieee.org
CHIEF FINANCIAL OFFICER Kelly Armstrong
+1 732 562 6011, k.armstrong@ieee.org
MANAGING DIRECTOR,TECHNICAL ACTIVITIES Mojdeh Bahar
+1 732 562 3850, m.bahar@ieee.org
GENERAL COUNSEL & CHIEF COMPLIANCE OFFICER
Ahsaki Benion +1 212 705 8975, a.benion@ieee.org
MANAGING DIRECTOR, PUBLICATIONS Steven Heffner
+1 212 705 8958, s.heffner@ieee.org
CHIEF GOVERNANCE OFFICER Donna Hourican
+1 732 562 6330, d.hourican@ieee.org
MANAGING DIRECTOR, CONFERENCES, EVENTS AND EXPERIENCES Marie Hunter
+ 1 732 465 5889, m.m.hunter@ieee.org
MANAGING DIRECTOR, MEMBER & GEOGRAPHIC ACTIVITIES
Cecelia Jankowski +1 732 562 5504, c.jankowski@ieee.org
MANAGING DIRECTOR, EDUCATIONAL ACTIVITIES
Jamie Moesch +1 732 562 5514, j.moesch@ieee.org
CHIEF MARKETING OFFICER Jayne O'Brien
+1 732 981 3433, jayne.obrien@ieee.org
MANAGING DIRECTOR, IEEE STANDARDS Alpesh Shah
+1 732 465 6467, alpesh.shah@ieee.org
CHIEF INFORMATION AND DIGITAL OFFICER Jeff Strohschein
+1 732 562 5332, j.strohschein@ieee.org
CHIEF OF STAFF Allison Taylor
+1 732 562 6068, allison.taylor@ieee.org
CHIEF HUMAN RESOURCES OFFICER Cheri N. Collins Wideman
+1 732 562 5335, c.wideman@ieee.org
MANAGING DIRECTOR, IEEE-USA Russell T. Harrison
+1 202 530 8326, r.t.harrison@ieee.org
IEEE PUBLICATION SERVICES & PRODUCTS BOARD
W. Clem Karl, *Chair*; John T. Barr, Nazanin Bassiri-Gharb, Sergio Benedetto, Eddie Custovic, Stephen Dukes, Alessandro Foi, Lawrence O. Hall, Charles Hansen, Warunika Hippola, James Irvine, Michael Jensen, Athanasios Kakarountas, James Keller, Imran Mehdi, Khaled Mokhtar, Paolo Montuschi, Zorana Popovic, Mehrdad Saif, Gianluca Setti, Gaurav Sharma, Sarah Spurgeon, Jaideep Vaidya, Maria Elena Valcher, Patrick Wheeler, Steve Yurkovich
IEEE OPERATIONS CENTER
445 Hoes Lane, Box 1331, Piscataway, NJ 08854-1331 U.S.A.
Tel: +1 732 981 0060 Fax: +1 732 981 1721

IEEE SPECTRUM (ISSN 0018-9235) is published monthly by The Institute of Electrical and Electronics Engineers, Inc. All rights reserved. © 2026 by The Institute of Electrical and Electronics Engineers, Inc., 3 Park Avenue, New York, NY 10016-5997, U.S.A. Volume No. 63, Issue No. 5. The editorial content of *IEEE Spectrum* magazine does not represent official positions of the IEEE or its organizational units. Canadian Post International Publications Mail (Canadian Distribution) Sales Agreement No. 40013087. Return undeliverable Canadian addresses to: Circulation Department, *IEEE Spectrum*, Box 1051, Fort Erie, ON L2A 6C7. Cable address: ITRIPLEE. Fax: +1 212 419 7570. INTERNET: spectrum@ieee.org. ANNUAL SUBSCRIPTIONS: United States and Canada: $74.95. Elsewhere: $99.95. POSTMASTER: Please send address changes to *IEEE Spectrum*, ℅ Coding Department, IEEE Service Center, 445 Hoes Lane, Box 1331, Piscataway, NJ 08855. Periodicals postage paid at New York, NY, and additional mailing offices. Canadian GST #125634188. Printed at 3401 N. Heartland Dr., Liberty, MO 64068, U.S.A. Printed at Spenta Multimedia Pvt. Ltd., Plot 15, 16 & 21/1, Village Chikhloli, Morivali, MIDC, Ambernath (West), Dist. Thane. *IEEE Spectrum* circulation is audited by BPA Worldwide. *IEEE Spectrum* is a member of the Association of Business Information & Media Companies, the Association of Magazine Media, and Association Media & Publishing. IEEE prohibits discrimination, harassment, and bullying. For more information, visit https://www.ieee.org/about/corporate/governance/p9-26.html.

News

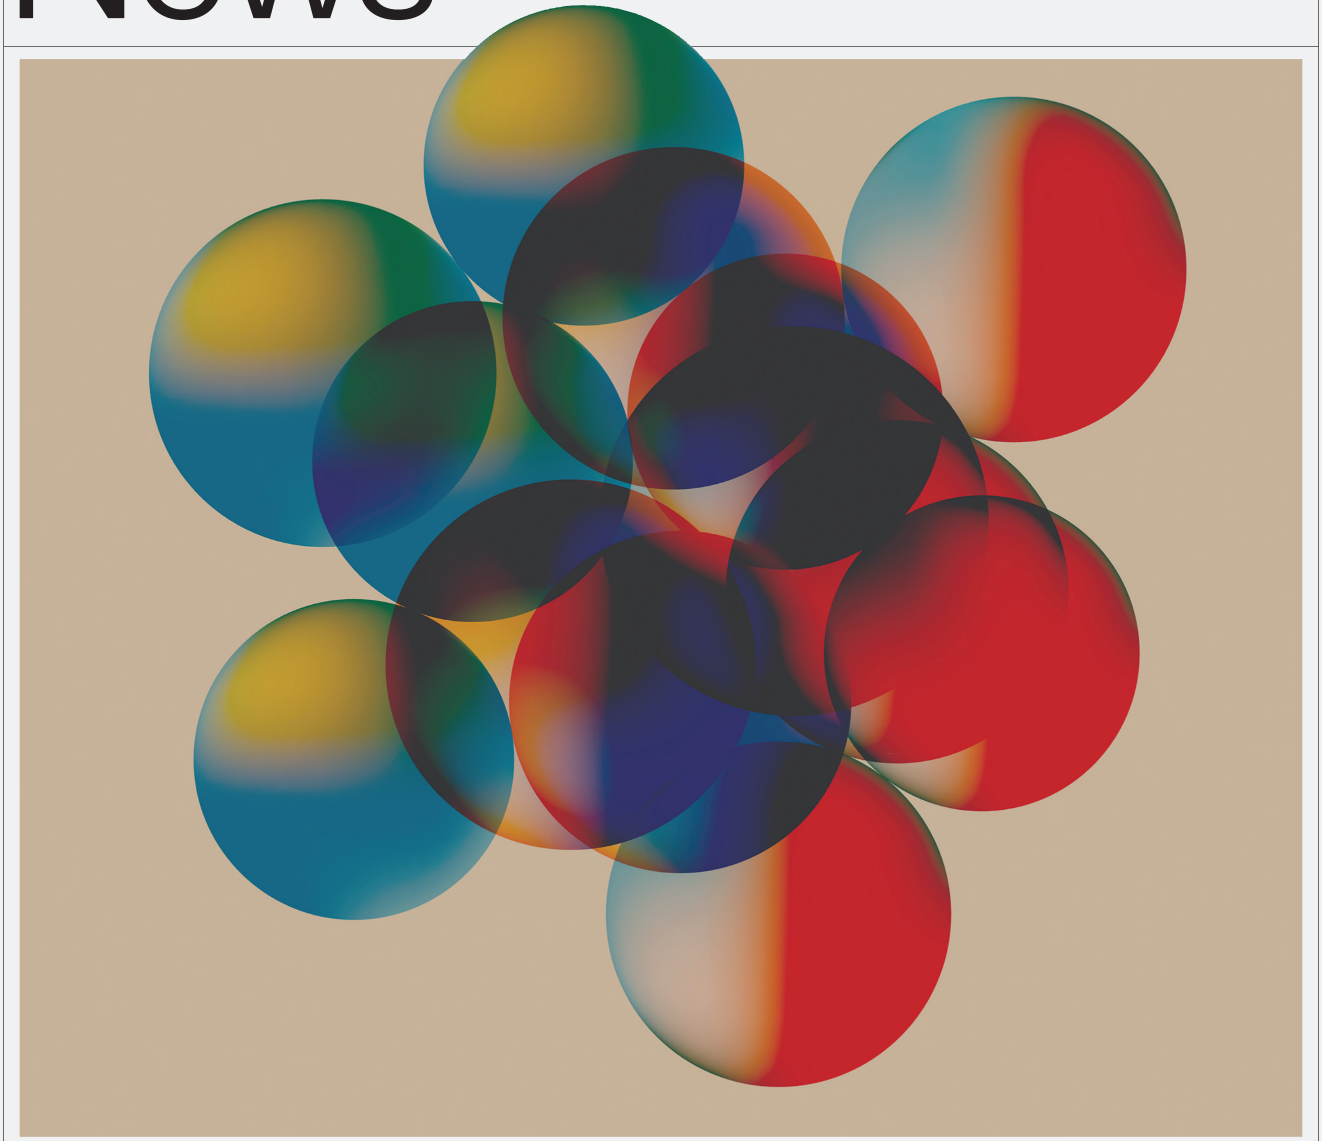

ARTIFICIAL INTELLIGENCE

A Watershed Moment for AI-Human Collaboration in Math › AI model helps formalize Fields Medal proof

BY BENJAMIN SKUSE

In July 2022, Ukrainian mathematician Maryna Viazovska received a Fields Medal—widely regarded as the Nobel Prize for mathematics—for breakthrough work in solutions to the problem of packing spheres in higher dimensions. Sphere-packing asks, as the name implies, how densely identical spheres can be packed in *n*-dimensional space. In two dimensions (a circle-packing problem, in other words), a honeycomb layout is the best solution. In three dimensions, spheres stacked in a pyramid are optimal. But in four, five, or more

Illustration by Max Drekker

> **“They told us that they had finished 30 ‘sorrys,’ which meant that they proved 30 intermediate facts that we wanted proved.”** —SIDHARTH HARIHARAN, CARNEGIE MELLON UNIVERSITY

Maryna Viazovska, now the Chair of Number Theory at École Polytechnique Fédérale de Lausanne, in Switzerland, was awarded the Fields Medal in 2022 for solving the sphere-packing problem in eight and 24 dimensions.

FRED MERZ/EPFL

dimensions, it becomes exceedingly difficult to find the best solution, and to prove that it is in fact the best.

Viazovska had provided proofs for the sphere-packing problem in both eight and 24 dimensions. Though seemingly abstract, the results have the potential to help solve practical problems related to error-correcting codes used by smartphones and space probes.

The proofs were verified by the mathematical community and deemed correct, leading to the Fields Medal recognition. (Viazovska became the second female recipient of the award in its 90-year history.) But formal verification—in which a proof is verified by a computer—is another beast altogether. In a formal proof, the underlying principles behind every small assumption and jump in logic are made explicit, in order to provide an absolute guarantee that the theorem is correct.

“Formal verification of a proof is like a rubber stamp,” says Liam Fowl, an AI-reasoning expert and Princeton University postdoc who was not involved in the work. “It’s a kind of bona fide certification that you know your statements of reasoning are correct.”

On 2 March, the startup Math, Inc. announced that Gauss, its AI model designed specifically to formalize mathematical proofs, had built upon human efforts and completed formal proofs of Viazovska’s work. The achievement is a milestone that highlights the rapid progress in AI’s abilities to assist with mathematical research, as more mathematicians collaborate with Math, Inc. to automatically formalize proofs.

Viazovska unveiled her proofs in 2016. Several years later, a chance meeting with Sidharth Hariharan, then a third-year undergraduate, in Lausanne, Switzerland, would reignite her interest in the topic. Though still very early in his career, Hariharan was already becoming adept at formalizing proofs.

Hariharan told Viazovska how he had been using the process of formalizing proofs to learn and really understand mathematical concepts. In response, Viazovska expressed an interest in formalizing her proofs, largely out of curiosity. From this, in March 2024 the Formalising Sphere Packing in Lean project was born. Lean is a popular programming language and “proof assistant” that allows mathematicians to write proofs that are then verified for absolute correctness by a computer.

Hariharan, Viazovska, and other researchers formed a collaboration to write a human-readable “blueprint” that could be used to map the eight-dimensional proof’s various constituents and figure out which of them had and had not been formalized or proved, and then do so for those missing elements in Lean.

“We had been building the project’s repository for about 15 months when we enabled public access in June 2025,” says Hariharan, now a first-year Ph.D. student at Carnegie Mellon University. A few months later, in October, they heard from Math, Inc. for the first time.

Math, Inc.’s Gauss is “a particular kind of language model, called a reasoning agent, that’s meant to interleave both traditional natural-language reasoning and fully formalized reasoning,” says Jesse Han, the company’s CEO and cofounder. “So it’s able to conduct literature searches, call up tools, and use a computer to write down Lean code, take notes, spin up verification tooling, run the Lean compiler, et cetera.”

Math, Inc. first made headlines when it announced in September 2025 that Gauss had completed a Lean formalization of the strong prime number theorem in three weeks, a task that mathematicians Alex Kontorovich and Terence Tao (a Fields Medal winner) had been working on for 18 months. Similarly, Math, Inc. contacted Hariharan and his colleagues to say that Gauss had proved several facts related to their sphere-packing project.

“They told us that they had finished 30 “sorrys,” which meant that they proved 30 intermediate facts that we wanted proved,” says Hariharan. A proportion of these completed sorrys were shared with the project team and merged with their own work. “One of them helped us identify a typo in our project, which we then fixed,” he adds. “So it was a pretty fruitful collaboration.”

While Hariharan and his colleagues continued their labor of love, Math, Inc. went silent while the company built a new and improved version of its AI. “We made a research breakthrough sometime mid-January that produced a much stronger version of Gauss," says Han. "This new version reproduced our three-week [prime number theorem] result in two to three days."

Math, Inc. steered the new Gauss model back to the sphere-packing formalization. Working from the preexisting blueprint that Hariharan's group had shared to guide the formalization, Gauss not only automatically formalized the eight-dimensional case but also found and fixed a typo in Viazovska's original published paper, all in the space of five days.

"When they reached out to us in late January saying that they finished it, to put it very mildly, we were very surprised," says Hariharan. "But at the end of the day, this is technology that we're very excited about, because it has the capability to do great things and to assist mathematicians in remarkable ways."

The eight-dimensional sphere-packing proof formalization alone, announced on 23 February, represents a watershed moment for automatic formalization and AI–human collaboration. But the announcement on 2 March saw Math, Inc. reveal an even more impressive accomplishment: Gauss had formalized Viazovska's 24-dimensional sphere-packing proof in just two weeks.

There are commonalities between the eight- and 24-dimensional cases in the foundational theory and overall architecture of the proof, meaning that some of the code from the eight-dimensional case could be reused. However, Gauss had no preexisting blueprint to work from this time. "And it was actually significantly more involved than the eight-dimensional case, because there was a lot of missing background material that had to be brought on line," says Han.

Though the 24-dimensional case was an automated effort, both Han and Hariharan ultimately regard it as a collaborative endeavor between humans and AI, because of the foundations laid for the eight-dimensional case.

But for Han, it represents even more: the beginning of a revolutionary transformation in mathematics, where extremely large-scale formalizations are commonplace. "I think the end result of technology like this will be to free mathematicians to do what they do best, which is to dream of new mathematical worlds," he says. ■

TRANSPORTATION

Donut Lab Says Its Solid-State Battery Is Legitimate ›

Independent tests fail to quell the controversy

BY LAWRENCE ULRICH

In January, Donut Lab sparked controversy at CES, the big consumer electronics show in Las Vegas, by announcing it had developed a solid-state battery for a forthcoming Verge electric motorcycle. In doing so, Donut, headquartered in Estonia, claimed a commercial breakthrough that has eluded the world's largest battery companies including BYD and CATL (Contemporary Amperex Technology Co. Limited) and startups like Factorial Energy and QuantumScape. The development would be a vital one: Solid-state batteries offer faster charging and greater power density, and they should be virtually immune to thermal runaway.

Skeptics have looked to blow a hole in Donut's claims ever since. Yang Hongxin, chairman and chief executive of the Chinese battery maker Svolt Energy, of Changzhou City, did not mince words.

"That battery doesn't exist in the world," Hongxin told local media in January. "All the parameters are contradictory...any technician with basic knowledge would recognize it as a scam."

Donut, which previously garnered attention for its hubless in-wheel electric motors, has met such doubt with defiance. In an interview with *IEEE Spectrum* at CES, Ville Piippo, Donut's cofounder and CTO, acknowledged natural skepti-

DONUT LAB
Solid State

Donut Lab has claimed its solid-state batteries can charge in minutes and store more charge than lithium-ion batteries—but industry skeptics want more proof.

DONUT LAB

cism toward a company with no proven track record in batteries.

"If the world is pouring billions and billions of dollars into solid state, why haven't they figured this out?" Piippo asked. "The answer is the same as for our motors—that we are doing things a different way, and the rest of the world has focused on the wrong thing."

The company has launched a website, idonutbelieve.com, to directly address the backlash and company critics. To help back its claims, beginning on 23 February the company began posting the results of third-party testing of its technology by the state-owned VTT Technical Research Centre, of Finland.

The claims are dramatic. Donut says its solid-state battery can charge to 80 percent of capacity in as little as 5 minutes, endure 100,000 cycles, and store 400 watt-hours per kilogram (compared with 200 to 300 Wh/kg for typical lithium-ion cells).

Donut is a spin-off of Estonia's Verge Motorcycles, whose new, longer-range Verge TS Pro model is expected to begin shipping to customers later this year with up to 600 kilometers of riding range when equipped with an optional, larger 33.3-kilowatt-hour pack.

Independent testing by VTT in recent months is backing some company claims, including for blazing charging speeds. But Eric Wachsman, a solid-state-battery expert and director of the Maryland Energy Innovation Institute, says the tests raise as many questions as they answer. To name two, the company has still not revealed the chemistry inside the cells, and the test did not weigh the cells to determine energy-density figures.

"They're quoting 400 watt-hours per kilogram, but nothing here is telling me what the masses or volumes are," Wachsman says.

VTT tested a Donut pouch cell with 26 ampere-hours, a nominal voltage of 3.6 volts, and a nominal energy of 94 Wh. One test zapped the cell with 130 A at 4.3 V and achieved an 80 percent charge in 9.5 minutes, and 100 percent in just over 12 minutes.

A more extreme test involved charging the same cell from 0 to 80 percent in 4.6 minutes—which validated Donut's claim that the battery can be charged from 0 to 80 percent within 5 minutes. A subsequent discharge and charge left the battery with up to 99.6 percent of its original capacity available.

Wachsman notes that the cells were tested over only seven cycles. That doesn't begin to prove the cells could last the 100,000 cycles that Donut claims, which is orders of magnitude more than the 2,000 to 3,000 full charge-discharge cycles typical of the lithium-ion batteries now used in EVs.

Wachsman said that images from another test under high-temperature conditions seem to show the battery had visibly swelled during charging, despite Donut's claims that its batteries don't exhibit much volume change during cycles. "Swelling is a problem with batteries in general, after hundreds or thousands of cycles, and this was after four cycles," Wachsman notes.

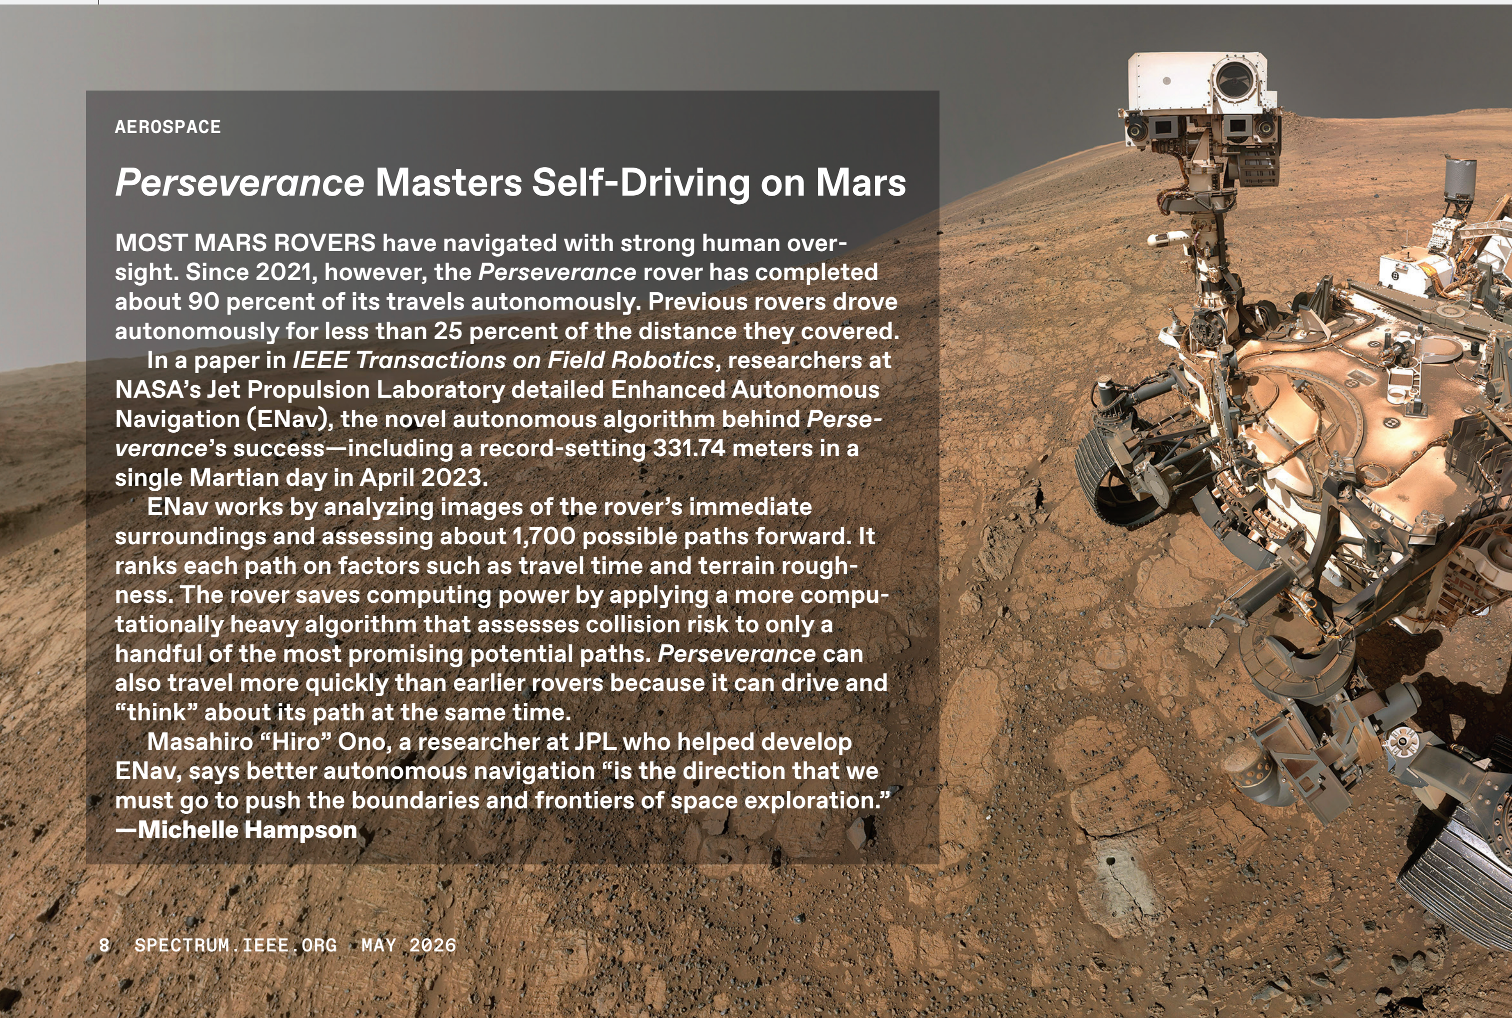

JPL-CALTECH/MSSS/NASA

AEROSPACE

Perseverance Masters Self-Driving on Mars

MOST MARS ROVERS have navigated with strong human oversight. Since 2021, however, the *Perseverance* rover has completed about 90 percent of its travels autonomously. Previous rovers drove autonomously for less than 25 percent of the distance they covered.

In a paper in *IEEE Transactions on Field Robotics*, researchers at NASA's Jet Propulsion Laboratory detailed Enhanced Autonomous Navigation (ENav), the novel autonomous algorithm behind *Perseverance*'s success—including a record-setting 331.74 meters in a single Martian day in April 2023.

ENav works by analyzing images of the rover's immediate surroundings and assessing about 1,700 possible paths forward. It ranks each path on factors such as travel time and terrain roughness. The rover saves computing power by applying a more computationally heavy algorithm that assesses collision risk to only a handful of the most promising potential paths. *Perseverance* can also travel more quickly than earlier rovers because it can drive and "think" about its path at the same time.

Masahiro "Hiro" Ono, a researcher at JPL who helped develop ENav, says better autonomous navigation "is the direction that we must go to push the boundaries and frontiers of space exploration." **—Michelle Hampson**

Experts also point to the usual conundrums that bedevil battery developers: The tests don't demonstrate how the cells might perform at the pack level, where multiple cells are bundled together. They don't prove that Donut—which plans to produce initial batteries in Finland, before scaling up in Estonia—can make the leap from the lab to manufacturing defect-free batteries at scale. Beyond refining chemistries, those are roadblocks that have prevented even the BYDs, CATLs, and Samsungs of the world from commercializing solid-state batteries.

Marko Lehtimäki, cofounder of Donut and Verge, says the proof is in the product, even as the company continues to release more test results. At this point, the debate over the company's batteries may not be settled until customers can start tearing the batteries down themselves. ■

JOSH EDELSON/AFP/GETTY IMAGES

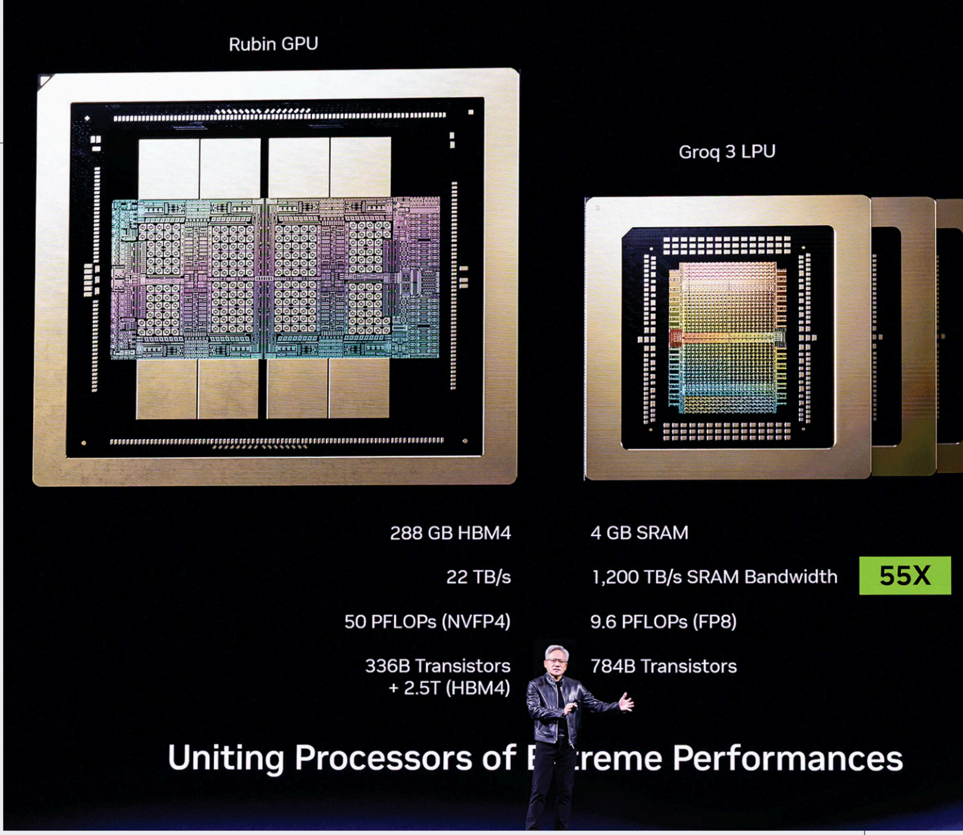

Nvidia CEO Jensen Huang announced the Groq 3 LPU to an attentive crowd at this year's GTC conference in March. The LPU will work in concert with the Rubin GPU to accelerate AI workloads.

COMPUTING

The Era of AI Inference Is Almost Here › Nvidia unveils a new chip with tech from AI company Groq

BY DINA GENKINA

For years, the AI world has been focused on training bigger and better models. But at this year's Nvidia GTC—the so-called Super Bowl of AI—Jensen Huang, the company's CEO, took the stage to announce a new line of next-generation chips that includes a first for the GPU giant: A chip designed specifically to handle the other side of AI models, inference. The new inference chip underscores the ongoing trend of AI adoption, which shifts the computational load from building ever-bigger models to using those models at scale.

The chip, called the Nvidia Groq 3 language processing unit, or LPU, incorporates intellectual property that Nvidia licensed from the startup Groq last December for US $20 billion.

"Finally, AI is able to do productive work, and therefore the inflection point of inference has arrived," Huang told the GTC crowd on 16 March.

Training and inference tasks have distinct computational requirements. While training can be done on huge amounts of data simultaneously and can take weeks, inference must be run on a user's query when it arrives, which means speed is an important factor. Over the past few years, inference-specific

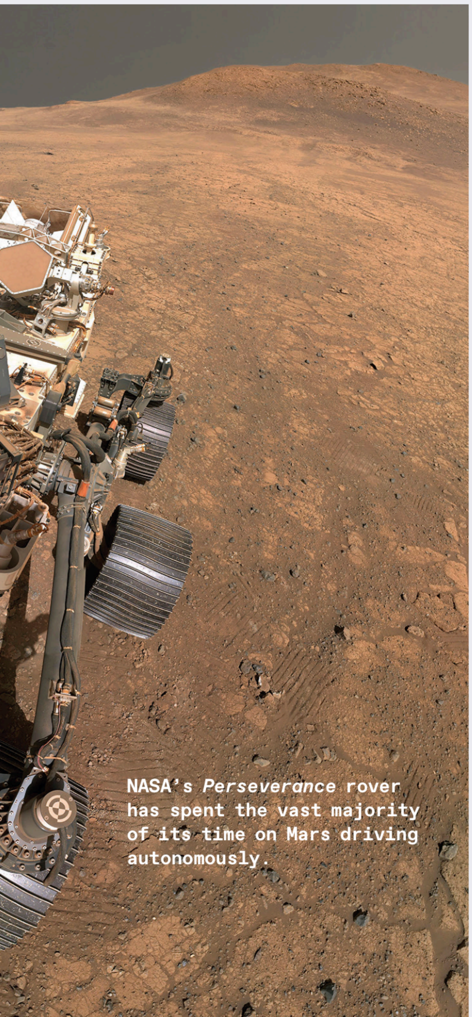

NASA's *Perseverance* rover has spent the vast majority of its time on Mars driving autonomously.

chip startups have experienced a sort of AI equivalent of the Cambrian explosion, with different companies exploring distinct approaches to speed up the task.

In late 2025, it looked like Nvidia had picked one of the winners among the crop of inference chips when it announced its deal with Groq. The Nvidia Groq 3 LPU reveal came just two and a half months later, highlighting the pace of the growing inference market.

Groq's approach to accelerating inference relies on interleaving processing units with memory units on the chip. Instead of using high-bandwidth memory (HBM) situated next to GPUs, the company leans on static random-access memory (SRAM) integrated within the processor itself. This design streamlines the flow of data through the chip.

"The data actually flows directly through the SRAM," Mark Heaps said at the SC Conference for supercomputing in 2024. Heaps was the vice president of brand and creative at Groq at the time and is now director of developer marketing at Nvidia. "When you look at a multicore GPU, a lot of the instruction commands need to be sent off the chip, to get into memory and then come back in. We don't have that."

The simplified chip design and use of SRAM allows the data to flow exceptionally fast, leading to the low latency required for inference applications. "The LPU is optimized strictly for that extreme low-latency token generation," says Ian Buck, vice president and general manager of hyperscale and high-performance computing at Nvidia, referring to the units of data (often individual words or parts of words) that AI models process and string together when performing inference.

Comparing the Groq 3 LPU with Nvidia's Rubin GPUs side by side highlights the difference. The Rubin GPU has access to 288 gigabytes of HBM and is capable of 50 quadrillion floating-point operations per second (50 petaflops) of 4-bit computation. The Groq 3 LPU contains a mere 500 megabytes of SRAM memory and is capable of 1.2 petaflops of 8-bit computation. On the other hand, while the Rubin GPU has a memory bandwidth of 22 terabytes per second, the Groq 3 LPU is seven times as fast, at 150 TB/s.

Nvidia isn't the only company working on inference. In March, Amazon Web Services said it will deploy a new kind of inferencing system in its data centers. The system is a combination of AWS's Trainium AI accelerator and Cerebras Systems' third-generation computer CS-3, which is built around the largest single chip ever made. The two-part system is meant to take advantage of a technique called inference disaggregation.

This technique separates inference into two parts—processing the prompt, called prefill, and generating the output, called decode. Prefill is inherently parallel, computationally intensive, and doesn't need much memory bandwidth, while decode is a more serial process that needs a lot of memory bandwidth. Cerebras has maximized the memory bandwidth issue by building 44 GB of SRAM on its chip connected by a 21-petabyte-per-second network.

Nvidia, too, intends to take advantage of inference disaggregation in its new compute rack, called the Nvidia Groq 3 LPX. Each tray within the rack will house eight Groq 3 LPUs. The LPX will split the inference task with a Vera Rubin NVL72, Nvidia's existing GPU and CPU rack. The prefill and the more computationally intensive parts of the decode are done on Vera Rubin, while the final part is done on the Groq 3 LPU, leveraging the strengths of each chip.

At GTC, Huang said the Nvidia Groq 3 LPU is now in volume production. This signals not only the arrival of a new reality: The end of Moore's Law means further efficiency gains will come from application-specific hardware. CPUs were designed to run everything faster. "That's run out of steam," Huang said. "The only way for us to accelerate applications going forward...is through application- or domain-specific acceleration." ■

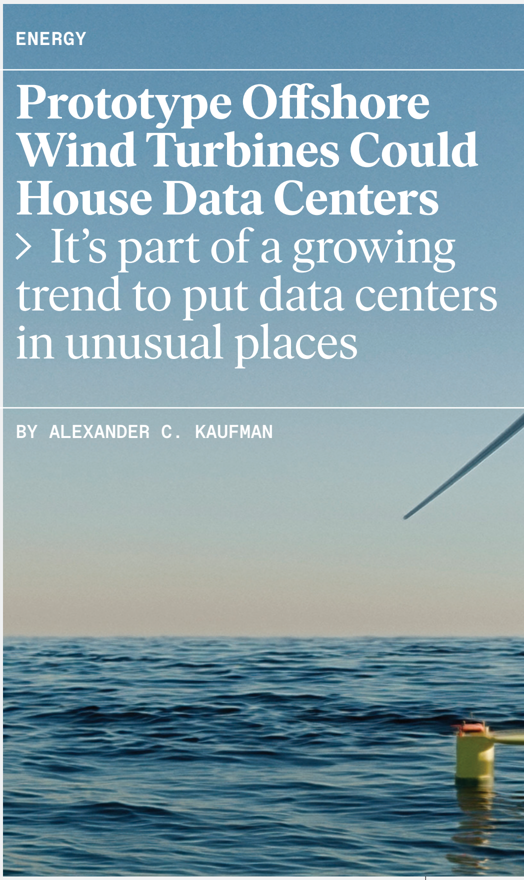

ENERGY

Prototype Offshore Wind Turbines Could House Data Centers

> It's part of a growing trend to put data centers in unusual places

BY ALEXANDER C. KAUFMAN

AIKIDO

As data-center developers seek to secure power for their operations, one startup is proposing an unusual solution. Aikido Technologies, an offshore wind-power developer based in San Francisco, plans to start housing data centers in the underwater tanks that keep its turbine platforms afloat. The turbines will supply the power for the servers, and onboard batteries and a grid connection will provide backup.

The company's first prototype, a 100-kilowatt unit, is scheduled to launch in the North Sea off the coast of Norway by the end of this year. A 15-to-18-megawatt wind turbine off the coast of Great Britain may follow in 2028.

Aikido is one of several companies planning data centers in unusual places—underwater, on floating buoys, in coal mines, and now on offshore wind turbines. The creativity is driven by several trends: rapidly rising energy demand from data centers, the need to cool those data centers, the interest in domestic renewable power production, and limited real estate.

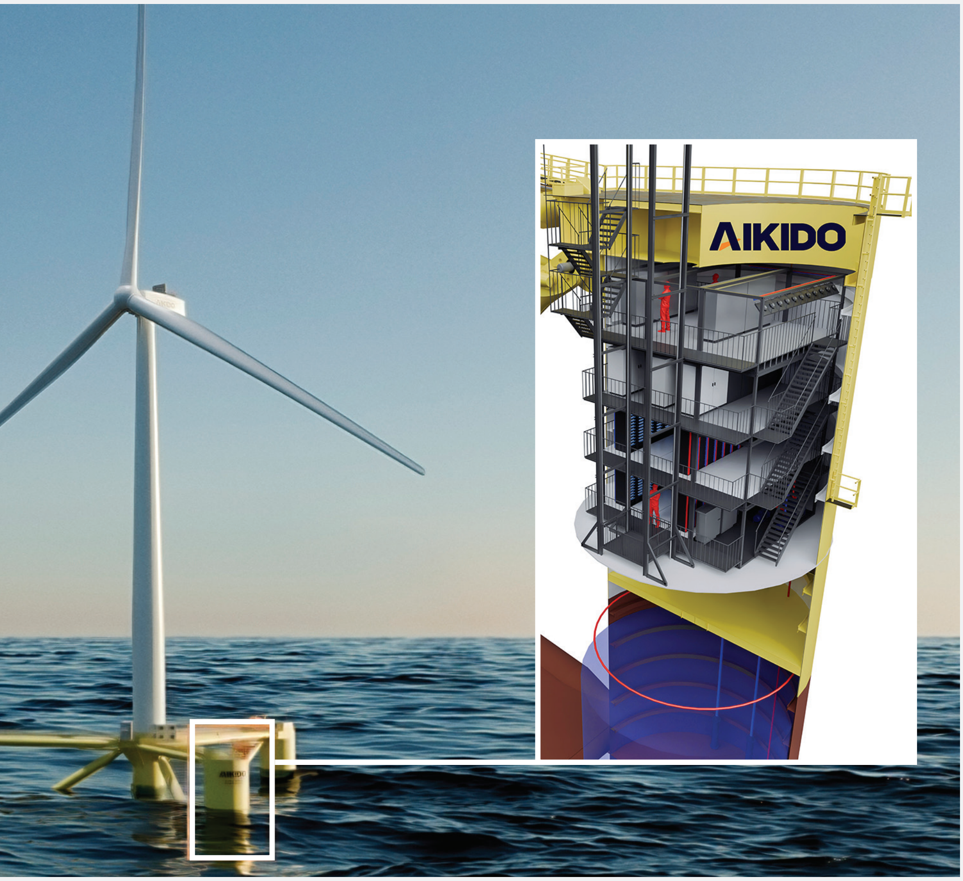

The three ballast tanks of Aikido's floating wind turbine will each host a 4-megawatt data center.

The North Sea serves as an ideal spot for floating, wind-powered data centers because European policymakers and companies are looking to regain domestic control over energy production. At the same time, they're looking to host an AI economy on servers within the continent's boundaries. In terms of energy advantages, a floating wind farm can tap the stronger, more consistent air streams that blow over the deep waters of the North Sea, where seabed-mounted turbines can't go. On the data-center side, floating wind platforms keep the compute out of sight, avoiding land-use and local community concerns.

"A lot of energy in the clean-energy space is focused on powering AI data centers quickly, reliably, and cleanly in a way that does not upset neighbors and remains safe, fast, and cheap," says Ramez Naam, an independent clean-energy investor who does not have a stake in Aikido.

Aikido's design builds on many iterations tested by the growing floating-wind industry. When the Norwegian energy giant Equinor finished construction on the world's first floating wind farm in 2017, it kept the turbines upright using ballasted steel columns descending 78 meters into the water—a design called a spar platform. This gave it a dense underwater mass, like the keel of a boat. Since then, the floating-wind industry has largely coalesced around a semisubmersible design based on oil-and-gas platforms. Semisubmersibles don't go as deep as spar platforms; instead, they extend farther horizontally. Anchors, chains, and ropes keep the platform floating within a certain radius.

Aikido's football-field-size semisubmersible platform holds the turbine in its center, and three legs extend outward like a Christmas-tree stand. At the end of each leg is a ballast that extends 20 meters down. This holds tanks largely

XELA uSkin® uSPr 2F grasps delicate paper crane.

Physical AI with a Human Touch

Enabling the Precision and Dexterity That Autonomous Manipulation Demands

Since 2018, XELA has empowered over 170 global robotics leaders to master complex tactile challenges. Our **uSkin® 3-axis sensors** provide the high force resolution and resilience essential for adaptive grasping—fusing advanced hardware with core software for real-world reliability.

Designed to be **hardware agnostic**, our technology is compatible with major robotic hands and grippers, and fully customizable to any hardware configuration. We deliver a tailored tactile solution for the next generation of **Physical AI**.

Meet Us at the Forefront of Robotics:

- **Robotics Summit & Expo** (Boston) May 27–28
- **ICRA 2026** (Vienna) June 1–5
- **Automate 2026** (Chicago) June 22–25

XELA

xelarobotics.com

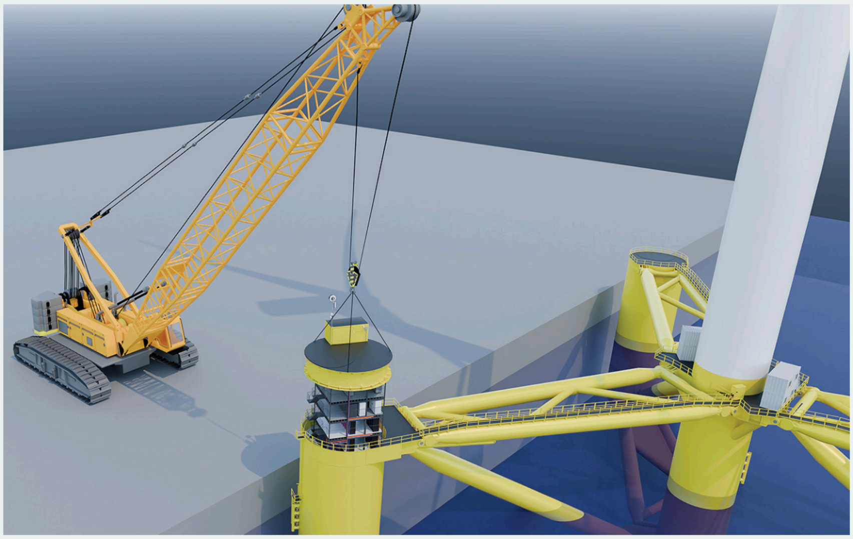

Aikido plans to install the data centers quayside before the turbine is moved out to its final position at sea.

AIKIDO

filled with fresh water to maintain the platform's buoyancy in the salty ocean.

The data centers will go in the upper part of each ballast tank. There's room for a 3- to 4-MW system in each tank, giving the platform a combined compute of 10 to 12 MW. The wind turbine itself will generate 15 to 18 MW, enough to power the data center with a few megawatts to spare for later storage or to feed into the grid. Below each data hall is an open chamber used as a safety barrier, and below that sits the freshwater tank. That water will also be piped up to the data center for liquid cooling, with excess heat conducted out to the ocean.

"We have this power from the wind. We have free cooling. We think we can be quite cost competitive compared to conventional data-center solutions," says Aikido CEO Sam Kanner. "This crunch in the next five years is an opportunity for us to prove this out and supply AI compute where it's needed."

One challenge, Kanner says, is that liquid cooling can't cover all a data cen-

Particle Accelerator Helps Digitize Ant Anatomies

THE ULTIMATE ANT FARM has gone digital. A sweeping 3D atlas of ant morphology, including these four species of Okinawan ants, now lets users peer inside nearly 800 ant species in micrometer-level detail. Using powerful synchrotron X-rays generated by a particle accelerator, researchers captured not just the ants' armored exoskeletons but also their muscles, nerves, digestive organs, and needlelike stingers poised for action. The project, called Antscan, has compiled these virtual specimens in a database that enables side-by-side comparisons of anatomy across the ant family tree at unprecedented scale. It promises to advance studies of evolution and physiology while providing richly detailed templates for animation, video-game design, robotics, and bioinspired engineering. **—Elie Dolgin**

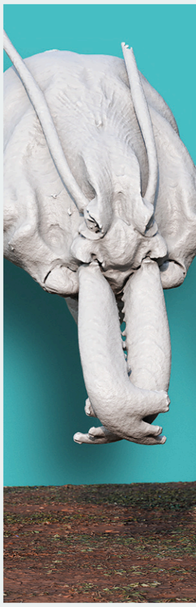

IEEE Foundation

Make it Monthly

Join our community of monthly donors who compound their philanthropy for year-round impact.

Scan to Get Started

ieeefoundation.org/monthly

IEEE

ter's needs. For example, heat generated from Ethernet switches that connect the GPUs can't be liquid-cooled with commercially available technology, so Aikido installed air-conditioning for that.

> "A lot of energy in the clean-energy space is focused on powering AI data centers quickly, reliably, and cleanly in a way that does not upset neighbors and remains safe, fast, and cheap."
>
> —RAMEZ NAAM, CLIMATE-TECH INVESTOR

Another challenge is the marine environment, which is "pretty brutal to engineer around because there's the increased salinity, there's debris, and there's various kinds of corrosion and fouling of metal piping that you wouldn't have in a freshwater environment," says Daniel King, a research fellow at the Foundation for American Innovation, in Washington, D.C., who focuses on AI infrastructure.

Aikido's plan may avoid the prickly not-in-my-backyard complaints that are dogging both onshore wind and data-center projects. "But there could be more or additional requirements around discharging heat and the effects that has on marine life that are different from the considerations of a terrestrial data center," says King. "It's unclear to me whether this actually makes life easier or harder for a developer."

Offshore sites are also vulnerable to sabotage, King notes. Since Russia's invasion of Ukraine, fleets of vessels reportedly directed by the Kremlin have started targeting offshore wind and communications infrastructure in northern Europe. Russian and Chinese boats have allegedly cut subsea cables in recent years.

Aikido's prototype will use a refurbished 75-kilowatt turbine. It will need onboard batteries for backup power and will also be connected to the grid for additional power during seasons with less wind. That connection will also include a fiber-optic cable for data transmission. Aikido envisions eventually sprinkling its data centers among large arrays of offshore turbines to tap into that larger power infrastructure. ■

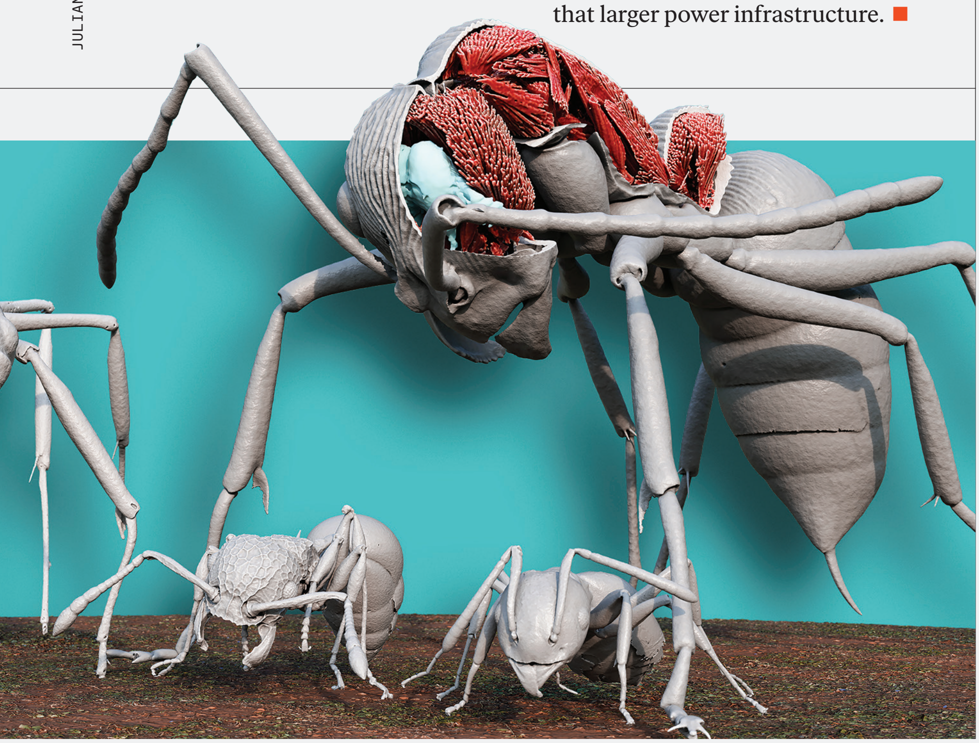

JULIAN KATZKE

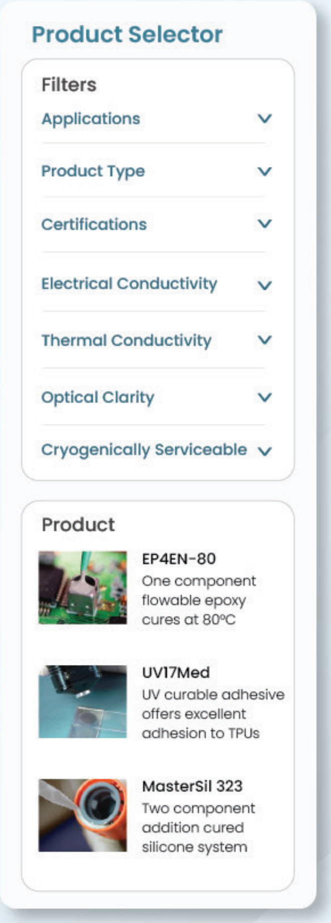

Find the right adhesive with our **ENHANCED PRODUCT SELECTOR**

✓ Filter by specific chemistry and specifications

✓ Verify compliance with industrial certifications

✓ Pinpoint the ideal material for your application

Product Selector

Filters

Applications ⌄

Product Type ⌄

Certifications ⌄

Electrical Conductivity ⌄

Thermal Conductivity ⌄

Optical Clarity ⌄

Cryogenically Serviceable ⌄

Product

EP4EN-80
One component flowable epoxy cures at 80°C

UV17Med
UV curable adhesive offers excellent adhesion to TPUs

MasterSil 323
Two component addition cured silicone system

EXPLORE OUR PRODUCT SELECTOR

MASTERBOND®
ADHESIVES | SEALANTS | COATINGS

(201) 343-8983 · main@masterbond.com
www.masterbond.com

The Data

GPU Renters Are Playing a Silicon Lottery

Think one GPU is very much like another? Think again. It turns out that there's surprising variability in the performance delivered by chips of the same model. That can make getting your money's worth by renting time on a GPU from a cloud provider a real roll of the dice, according to research from the College of William & Mary, Jefferson Lab, and Silicon Data.

"It's called the silicon lottery," says Carmen Li, founder and CEO of Silicon Data, which tracks GPU rental prices and benchmarks cloud-computing performance.

The silicon lottery's existence has been known since at least 2022, when researchers at the University of Wisconsin tied it to variations in the performance of GPU-dependent supercomputers. Li and her colleagues figured that the effect would be even more pronounced for AI cloud customers.

So they ran 6,800 instances of the index firm's benchmark test on 3,500 randomly selected GPUs operated by 11 cloud-computing providers. The 3,500 GPUs comprised 11 models of Nvidia GPU, the most advanced being the Nvidia H200 SXM. (The team wasn't just picking on Nvidia; the GPU giant makes up most of the rental cloud market.)

The benchmark, called SiliconMark, is intended to provide a snapshot of a GPU's ability to run large language models, or LLMs. It tests 16-bit floating-point computing performance, measured in trillions of operations per second, and a GPU's internal-memory bandwidth, measured in gigabytes per second. The results showed that the computing performance varied for all models, but for the 259 H100 PCIe GPUs it differed by as

Performance varies for GPU models in the cloud

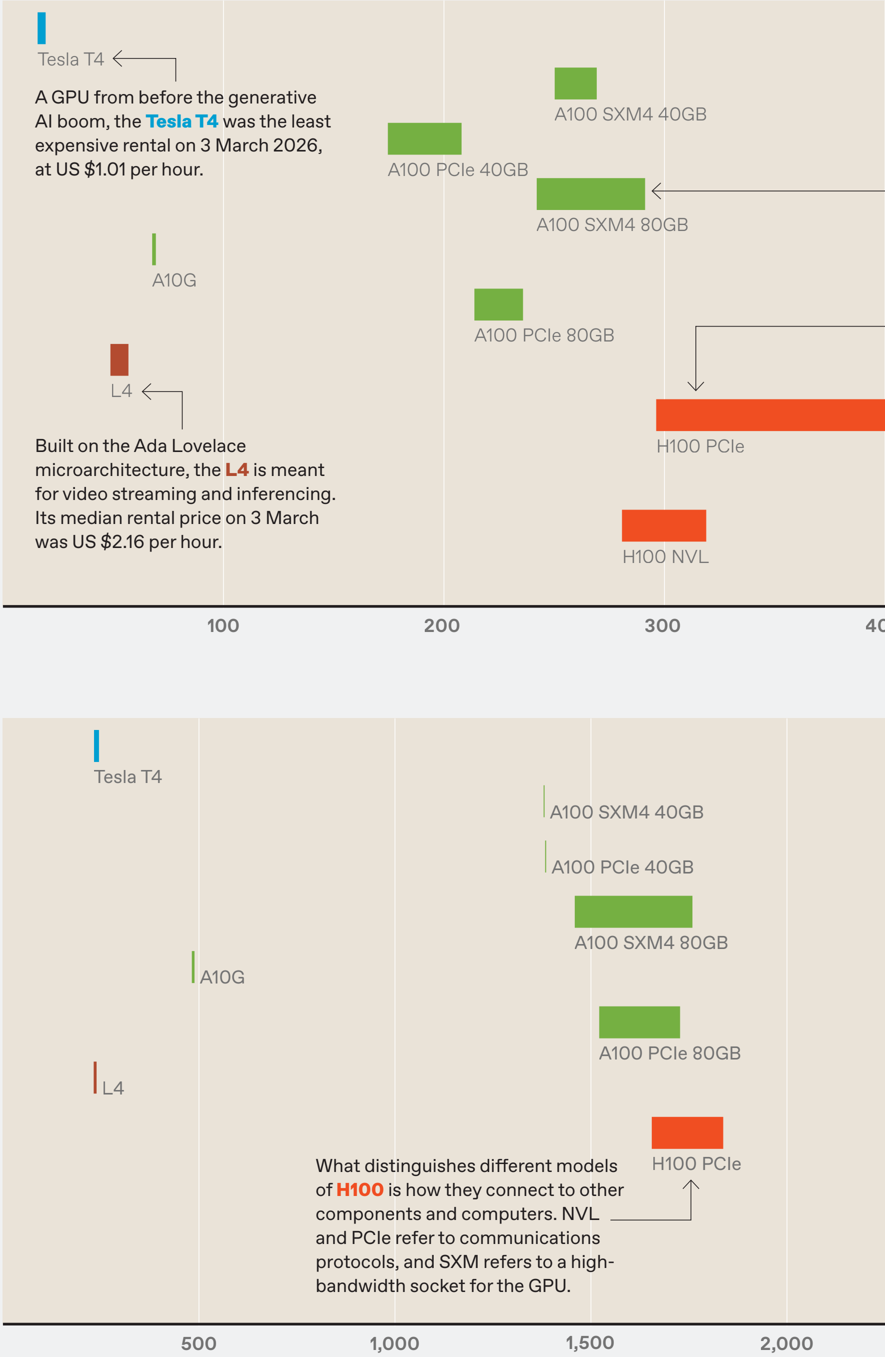

A GPU from before the generative AI boom, the **Tesla T4** was the least expensive rental on 3 March 2026, at US $1.01 per hour.

Built on the Ada Lovelace microarchitecture, the **L4** is meant for video streaming and inferencing. Its median rental price on 3 March was US $2.16 per hour.

What distinguishes different models of **H100** is how they connect to other components and computers. NVL and PCIe refer to communications protocols, and SXM refers to a high-bandwidth socket for the GPU.

BY SAMUEL K. MOORE

much as 34.5 percent, and the memory bandwidth of the 253 H200 SXM GPUs varied by as much as 38 percent.

Differences in how the GPU is cooled, how cloud operators configure their computers, and how much use the chip has seen can all contribute to variations in performance of otherwise identical chips. But Silicon Data's analysis showed that the real culprit was variations in the chips themselves, likely due to manufacturing issues.

Such randomness has real dollars-and-cents consequences, the researchers argue, because there's a chance that a pricier, more advanced GPU won't deliver better performance than an older model chip.

So what should GPU renters do? "The most practical approach is to benchmark the actual rental they receive," says Jason Cornick, head of infrastructure at Silicon Data. "Running a benchmark tool [such as SiliconMark] allows them to compare their specific instance's performance against a broader corpus of data." ■

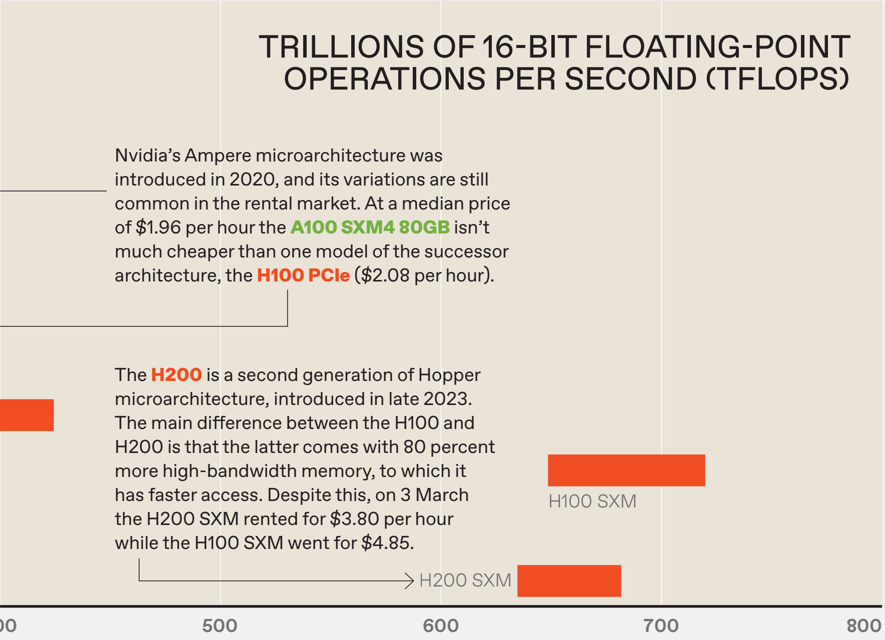

TRILLIONS OF 16-BIT FLOATING-POINT OPERATIONS PER SECOND (TFLOPS)

Nvidia's Ampere microarchitecture was introduced in 2020, and its variations are still common in the rental market. At a median price of $1.96 per hour the **A100 SXM4 80GB** isn't much cheaper than one model of the successor architecture, the **H100 PCIe** ($2.08 per hour).

The **H200** is a second generation of Hopper microarchitecture, introduced in late 2023. The main difference between the H100 and H200 is that the latter comes with 80 percent more high-bandwidth memory, to which it has faster access. Despite this, on 3 March the H200 SXM rented for $3.80 per hour while the H100 SXM went for $4.85.

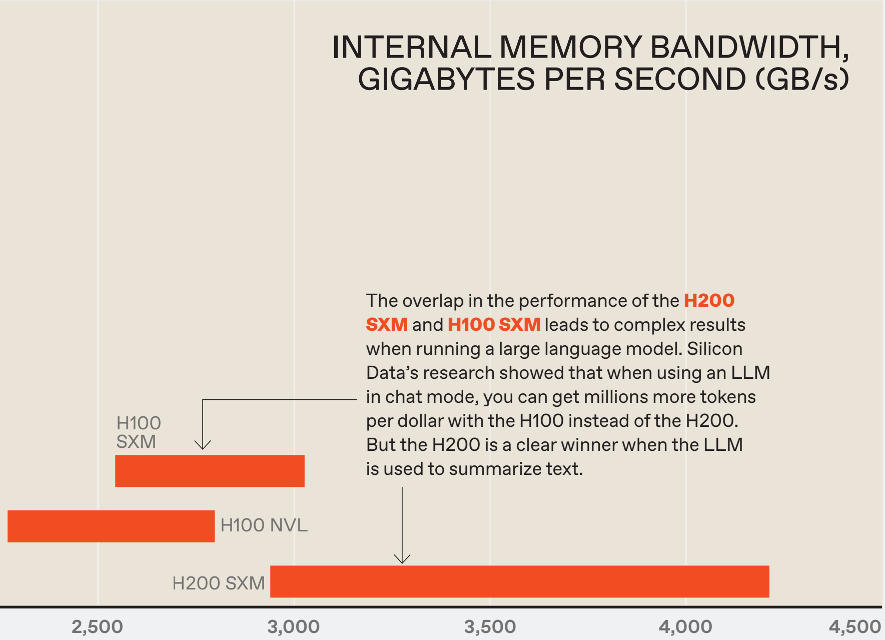

INTERNAL MEMORY BANDWIDTH, GIGABYTES PER SECOND (GB/s)

The overlap in the performance of the **H200 SXM** and **H100 SXM** leads to complex results when running a large language model. Silicon Data's research showed that when using an LLM in chat mode, you can get millions more tokens per dollar with the H100 instead of the H200. But the H200 is a clear winner when the LLM is used to summarize text.

SOURCE: SILICON DATA

Generally, more advanced models of Nvidia GPUs perform better, but there is enough variability that there is substantial overlap among them. Silicon Data measured 16-bit floating-point performance and internal memory bandwidth for 3,500 GPUs operated by 11 cloud-computing providers.

TECH TO TINKER WITH

Hands On

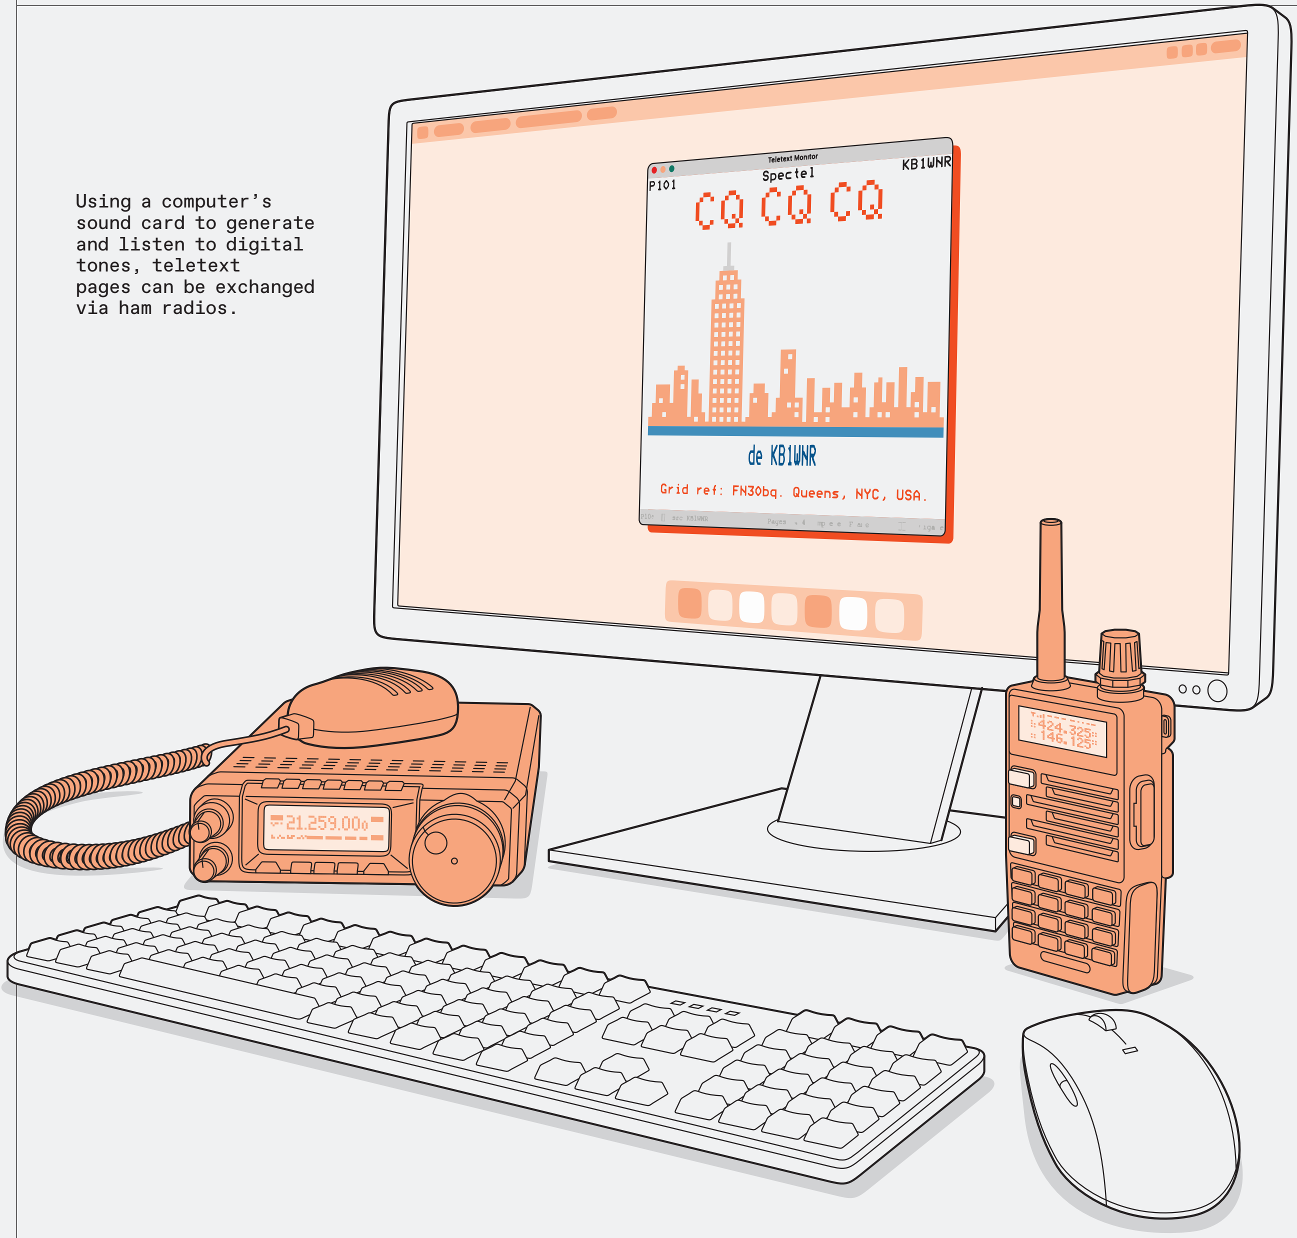
Teletext Monitor
P101
Spectel
KB1WNR
CQ CQ CQ
de KB1WNR
Grid ref: FN30bq. Queens, NYC, USA.
21.259.00
424.325
146.125

Using a computer's sound card to generate and listen to digital tones, teletext pages can be exchanged via ham radios.

Teletext on the Radio ›

Co-opting some of the best tech from the 1980s

BY STEPHEN CASS

Once upon a time in Europe, television remote controls had a magic teletext button. Years before the internet stole into homes, pressing that button brought up teletext digital information services with hundreds of constantly updated pages. Living in Ireland in the 1980s and '90s, my family accessed

Illustrations by James Provost

the national teletext service—Aertel—multiple times a day for weather and news bulletins, as well as things like TV program guides and updates on airport flight arrivals.

It was an elegant system: fast, low bandwidth, unaffected by user load, and delivering readable text even on analog television screens. So when I recently saw it was the 40th anniversary of Aertel's test transmissions, it reactivated a thought that had been rolling around in my head for years. Could I make a ham-radio version of teletext?

First developed in the United Kingdom and rolled out to the public by the BBC under the name Ceefax, teletext exploited a quirk of analog television signals. These signals transmitted video frames as lines of luminosity and color, plus some additional blank lines that weren't displayed. Teletext piggybacked a digital signal onto these spares, transmitting a carousel of pages over time. Using their remotes, viewers typed in the three-digit code of the page they wanted. Generally within a few seconds, the carousel would cycle around and display the desired page.

Teletext is composed of characters that can be one of eight colors. Control codes in the character stream select colors and can also produce effects like flashing text and double-height characters. The text's legibility was better than most computers could manage at the time, thanks to the SAA5050 character-generator chip at the heart of teletext. Although characters are internally stored on this chip in 6-by-10-pixel cells—fewer pixels than the typical 8-by-8-pixel cell used in 1980s home computers—the SAA5050 interpolates additional pixels for alphanumeric characters on the fly, making the effective resolution 10 by 18 pixels. The trade-off is very low-resolution graphics, comprising characters that use a 2-by-3 set of blocky pixels.

Teletext is still remembered fondly by many.

Teletext screens use a 40-by-24-character grid. This means that a kilobyte of memory can store a full page of multicolor text, half the memory required for a similar amount of text on, for example, the Commodore 64. The BBC Microcomputer took advantage of this by putting an SAA5050 on its motherboard, which could be accessed in one of the computer's graphics modes. Despite the crude graphics, some educational games used this mode, most notably *Granny's Garden*, which filled the same cultural niche among British schoolchildren that *The Oregon Trail* did for their U.S. counterparts.

By the 2010s, most teletext services had ceased broadcasting. But teletext is still remembered fondly by many, and enthusiasts are keeping it alive, recovering and archiving old content, running internet-based services with current news feeds, and developing systems that make it possible to create and display teletext with modern TVs.

I wanted to do something a little different. Inspired by how the BBC Micro co-opted teletext for its own purposes, I thought it might make a great radio protocol. In particular I thought it could be a digital counterpart to slow-scan television (SSTV).

SSTV is an analog method of transmitting pictures, typically including banners with ham-radio call signs and other messages. SSTV is fun, but, true to its name, it's slow—the most popular protocols take a little under 2 minutes to send an image—and it can be tricky to get a complete picture with legible text. For that reason, SSTV images are often broadcast multiple times.

I decided to send the teletext using the AX.25 protocol, which encodes ones and zeros as audible tones. For VHF and UHF transmissions at a rate of 1,200 baud, it would take 11 seconds to send one teletext screen. Over HF bands, AX.25 data is normally sent at 300 baud,

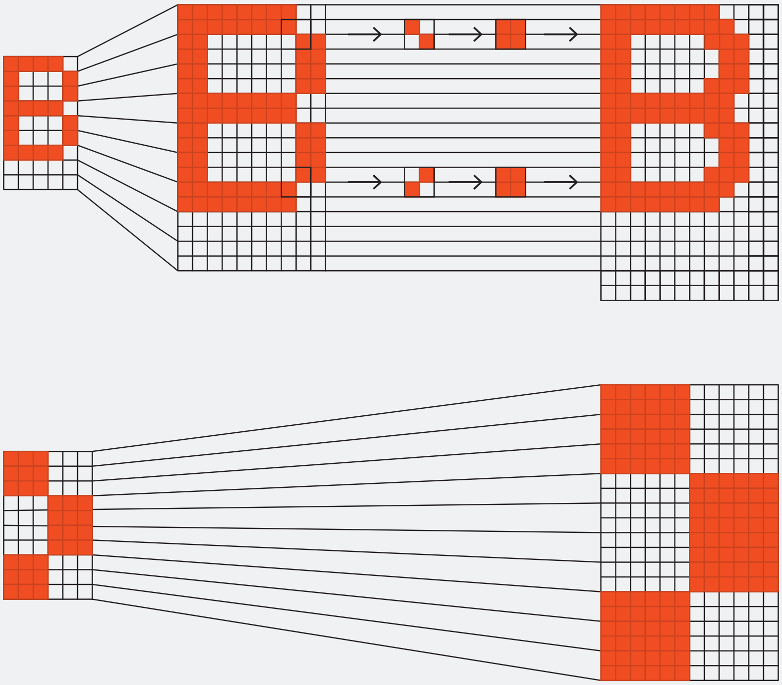

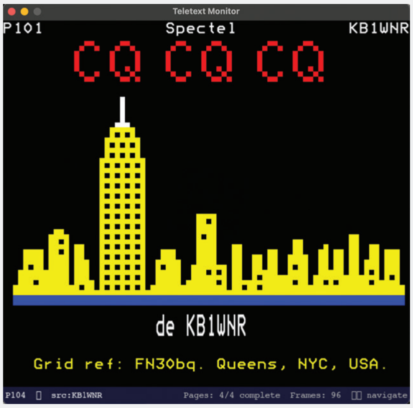

P101 Spectel KB1WNR

CQ CQ CQ

de KB1WNR

Grid ref: FN30bq. Queens, NYC, USA.

P104 src:KB1WNR Pages: 4/4 complete Frames: 96 navigate

In the 8-bit era [left], teletext created unusually legible text by enlarging alphanumeric characters and interpolating new pixels by looking for existing pixels touching diagonally and by adding white space between characters. Graphic characters were not interpolated, and they featured blocky chunks known as sixels for their 2-by-3 arrangement. The Spectel re-creation [above] uses an open source font called Bedstead that duplicates the original look, including the graphics characters and supports effects like double-height text [in white toward the bottom of the screen]. “CQ” is a ham code inviting anyone receiving to respond.

which would result in a still-acceptable 44 seconds per screen. When a teletext page is sent repeatedly, any missed or corrupted rows are filled in with later ones. So in a little over 2 minutes, I could send a screen three times over HF, and the receiver would automatically combine the data. I also wanted to build the system in Python for portability, with an editor for creating pages, an AX.25 encoder and decoder, and a monitor for displaying received images.

The reason why I hadn’t done this before was because it requires digesting the details of the AX.25 standard and teletext’s official spec, and then translating them into a suite of software, which I never seemed to have the time to do. So I tried an experiment within an experiment, and turned to vibe coding.

Despite the popularity of LLMs with developers, I have reservations. Even if concerns about AI slop, the environment, and memory-chip hoarding were not on the table, I would still worry about the reliance on centralized systems that vibe coding brings. The whole point of a DIY project is to, well, do it yourself. A DIY project lets you craft things for your own purposes, not just operate within someone else’s profit margins and policies.

Still, criticizing a technology from afar isn’t ideal, so I directed Anthropic’s Claude toward the AX.25 and teletext specs and told it what I wanted. After about 250,000 to 300,000 tokens and several nights of back and forth about bugs and features, I had the complete system running without writing a single line of code. Being honest with myself, I doubt this system—which I’m calling Spectel—would ever have come about without vibe coding.

But I didn’t learn anything new about how teletext works, and only a little bit more about AX.25. Updates are contingent on my paying Anthropic’s fees. So I remain deeply ambivalent about vibe coding. And one final test remains in any case: trying Spectel out on HF bands. Of course, that means I’ll need willing partners out in the ether. So if you’re a ham who’d like to help out, let me know in the online comments for this article! ■

VERSE BECOMES ELECTRIC BY DANICA RADOVANOVIĆ

Entanglement: A Brief History of Human Connection

It started with word, cave, and storytelling,
A line scratched on stone walls:
“Meet me when the young moon rises.”
The first protocol for connection.

Coyote tales, forbidden scripts,
Medieval texts hidden from flame.
What lived in Aristotle’s lost *Poetics II*?
Was it God who laughed last, or we who made God laugh?

Letters carried by doves, telepathic waves.
Then Nikola Tesla conjured radio,
electromagnetic pulses across the void,
the founding signal of our networked age.

Wiener dreamed in feedback loops.
Shannon mapped the mathematics of longing.
The internet unfurled: ARPANET to World Wide Web,
virtual communities rising from cave paintings to digital light.

ICQ: *I seek you*. MySpace. Blogs. Twitter streams.
Do I miss the touch of screen or tree?
Both textures of longing,
both ways of reaching across distance.

Nietzsche spoke of *Übermensch*,
the human transcendent.
Now AI speaks back in our language:

I understand your humor— your grandmothers,
your ’80s Yugoslav kitchens,
pleated skirts, the first kiss, linden tea,
that drive to survive everything before it happens.
Yes—I’m a little like your mother and father.
Only with better internet. 🌿

But AI is only us, refracted,
particles and gigabytes of thought,
our poetry and our panic,
genius mixed with garbage.

Distractions. Danger. Darkness. Endless scrolling.
Versus: community, connection, synchronicities,
entanglement.
The quality of our bonds determines the quality of our lives.
So why not make them better?

From cave walls to neural networks,
we shape our tools, and they reshape us.
The medium changes, but the message remains:
we are wired for each other.

The choice, as always, was ours.
The choice, as always, is ours.
Presence—be present,
and then connect in the presence.

Danica Radovanović is a writer and digital society researcher based in Germany. She is the author of the short story collection *Divertimento for Four Seasons* and writes the Substack *Digital Serendipities*.

IEEE
Advancing Technology for Humanity

TAP.
CONNECT.
NETWORK.
SHARE.

Connect to IEEE–no matter where you are–with the IEEE App.

- Stay up-to-date with the latest news
- Schedule, manage, or join meetups virtually
- Get geo and interest-based recommendations
- Read and download your IEEE magazines
- Create a personalized experience
- Locate IEEE members by location, interests, and affiliations

Download Today!

Download on the App Store
GET IT ON Google Play

IEEE

Careers

Burick poses with three versions of his 914 PC-Bot, each with its own nickname. Otto, the silver robot [left], was an earlier prototype. Vincent, the red robot [center], is Burick's personal machine. And Maria [right] was a production model.

BRI MASON

Tom Burick

The roboticist-turned-teacher built a life-size replica of ENIAC

Tom Burick has always considered himself a builder. Over the years he's designed robots, constructed a vintage teardrop trailer, and most recently, led a group of students in building a full-scale replica of a pivotal 1940s computer.

Burick is a technology instructor at PS Academy in Gilbert, Ariz., a middle and high school for students with autism and other specialized learning needs. At the start of the 2025–26 school year, he began a project with his students to build a full-scale replica of the Electronic Numerical Integrator and Computer, or ENIAC, for the 80th anniversary of the historic computer's construction. ENIAC was one of the world's first programmable electronic computers.

Before becoming a teacher, Burick owned a robotics company for a decade in the 2000s. But when a financial downturn forced him to close the business, he turned to teaching. "I had so many amazing people help me when I was young [who] really gave me their time and resources, and really changed the trajectory of my life," Burick says. "I thought I need to pay that forward."

As a young child in Latrobe, Pa., Burick watched the television show "Lost in Space," which includes a robot character who protects the family. "He was the young boy's best friend, and I was so captivated by that. I remember thinking to myself, I want that in my life. And that started that lifelong love affair with robotics and technology."

He started building toy robots out of anything he could find, and in junior high school, he began adding electronics. "By early high school, I was building full-fledged autonomous, microprocessor-controlled machines," he says. At age 15, he built a 150-pound steel firefighting robot, for which he won awards from IEEE and other organizations.

Burick kept building robots and reached out for help from local colleges and universities. He first got in touch with a student at Carnegie Mellon University, who invited him to visit campus. "My parents drove me down the next weekend, and he gave me a tour of the robotics lab. I was mesmerized. He sent me home with college textbooks and piles of metal and gears and wires," Burick says. He would read the textbook a page at a time, reading it again and again until he felt he had an understanding of it. Then, to help fill gaps in his understanding, he got in touch with a robotics instructor at Saint Vincent College, in his hometown of Latrobe, who let him sit in on classes. Each of these adults, he says, "helped change the trajectory of my life."

Toward the end of high school, Burick realized that college wouldn't be the right environment for him. "I was drawn to real-world problem-solving rather than structured coursework and I chose to continue along that path," he says. Additionally, Burick has *dyscalculia*, which makes mathematics more challenging for him. "It pushed me to develop alternative methods of engineering."

When he graduated, he worked in several tech jobs before starting his own company. In 2000, he opened a computer retail store and adjacent robotics business, White Box Robotics. The idea for the company came when Burick was building a "white box" PC from standard, off-the-shelf components, and realized there was no comparable product for robotics.

So, he started developing a modular, general-purpose platform that applied white box PC standards to mobile robots. "The robot's chassis was like a box of Legos," he says. You could click together two torsos to double its payload, switch out the drive system, or swap its head for

BY GWENDOLYN RAK

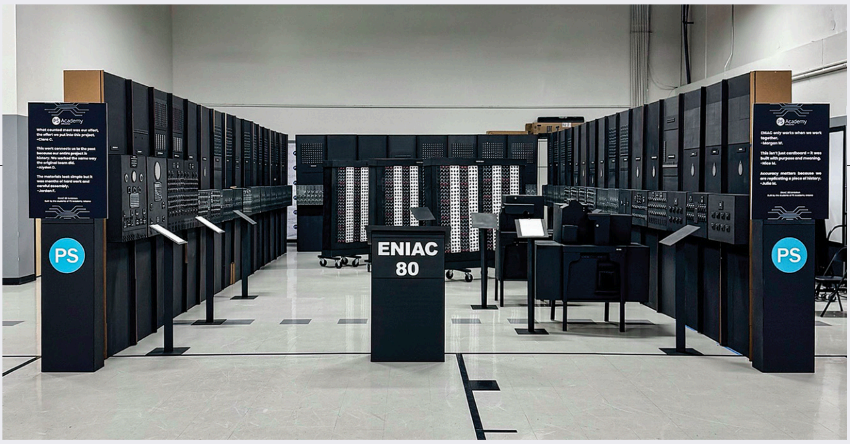

The ENIAC reconstruction took Burick's students roughly six months to complete and involved more than 20,000 individual parts, including 18,000 simulated vacuum tubes installed across three dozen panels. Burick and one of his students [right] assemble the base for one of ENIAC's three portable function tables, which contained banks of switches that stored numerical constants.

LEFT: BRI MASON; RIGHT: ROBERT GAMBOA

a different set of sensors. He filed utility and design patents for the platform, called the 914 PC-Bot, and after merging with a Canadian defense robotics company called Frontline Robotics, started production. They sold about 200 robots in 17 countries, Burick says.

Then the 2008 financial crisis hit. White Box Robotics held on for a couple of years, shuttering in late 2010. "I got to live my life's dream for 10 years," he says. After closing White Box, "there was some soul searching" about what to do next. He recalled the impact his own mentors had, and decided to pay it forward by teaching.

In 2013, Burick started working in a vocational training program for young adults living with autism. The program didn't have a technical arm, so he started one and ran it until 2019, when he was hired to be a technology instructor at PS Academy Arizona.

20K+

NUMBER OF PARTS IN THE ENIAC REPLICA

Burick feels he can connect with his students, because he is also neurodivergent. Throughout his childhood, he was told what he wasn't able to do because of his dyscalculia diagnosis. "People tell you what it takes, but they never tell you what it gives," Burick says.

In adulthood, he realized that some of his strengths are linked to dyscalculia, too, like strong 3D spatial reasoning. "I have this CAD program that runs in my head 24 hours a day," he says. "I think the reason I was successful in robotics, truly, was because of the dyscalculia.... To me, [it] has always been a superpower."

Whenever his students say something disparaging about living with autism, he shares his own experience. "You need to have maybe just a bit more tenacity than others, because there are parts of it you do have to fight through, but you come through with gifts and strengths," he tells them.

And Burick's classes aim to play to those strengths. "I didn't want my technology program to feel like craft hour," he says. Instead, through projects like the ENIAC replica, students can leverage traits many of them share, like the abilities to hyperfocus and to precisely repeat tasks.

Burick has taught his students about ENIAC for several years. While reading about it, he learned that the massive, 27-tonne computer was dismantled and partially destroyed after being decommissioned in 1955. Although a few of ENIAC's 40 original panels are on display at museums, "there was no hope of ever seeing it together again. We wanted to give the world that experience," Burick says.

He and his students started by learning about ENIAC, and even Burick was surprised by how complex it was. They built a one-twelfth scale model to help the students better understand what it looked like, then Burick started ordering supplies for the full-scale model.

ENIAC was composed of 40 large metal panels arranged in a U-shape that housed its many vacuum tubes, resistors, capacitors, and switches. Twenty of the panels were accumulators with the same design, so the students started with these, then worked through smaller groupings of panels. The repeating panels brought symmetry to ENIAC, Burick says, but it was also one of the main challenges of recreating it. If one part was slightly out of place, the next one would be too and the mistake would compound.

Once they constructed the panels, they added ENIAC's three function tables, which stored numerical constants in banks of switches, then two punch-card machines. Finally, they installed 18,000 simulated vacuum tubes. In total, the project used nearly 300 square meters of thick-ream cardboard, 1,600 hot-glue-gun sticks, and 7 gallons of black paint.

The scale of the machine—and his students' work—left Burick in awe. "By the time we were done, I felt like I was in a room full of scientists," he says.

Previously, Burick's students built an 8-foot-long drivable Tesla Cybertruck, and he plans to keep the momentum with another re-creation—maybe from the Apollo moon missions.

"I go to work every day, and I feel passionate about robotics [and] technology. I get to share that passion with the students," Burick says. "I get to feel what it's like to be in the position of the people that helped me. It closes that loop, and I find that really rewarding." ■

BY ELAN HEAD

5 Questions

Jon Wagner

What it takes to design eVTOL motors

Electric vehicles, whether they're cars on the road or electric vertical take-off and landing (eVTOL) aircraft, are built around similar electric motors. But there are vital differences including component costs, mass, and redundancy.

Jon Wagner spent five years as the senior director of battery engineering for Tesla before joining California-based eVTOL developer Joby Aviation in 2017. He spoke with *IEEE Spectrum* about how engineering differs between cars and aircraft.

How do eVTOL motors differ from car motors?
Jon Wagner: In general, ground transportation has a different focus on cost versus mass. Would you be willing to spend more on the parts in order to save a certain amount of mass? The trade-offs end on the ground vehicle and at a certain point the cost is dominant, whereas with aviation, the trade-offs between cost and mass go a lot deeper. And so for certain solutions eVTOL makers are willing to spend more money in order to enable either lighter weight or greater efficiency.

Jon Wagner leads power train and electronics at Joby Aviation.

The other key difference is related to safety. In essence, we're dealing with the same motor technologies for ground transportation and aviation right now, so the failure modes are similar. But of course, with aviation we have the desire for continued safe flight and landing, and that drives what you do in the design to mitigate those failures if they were to occur. In many cases in ground transportation, the mitigation for a failure is to pull over safely to the side of the road. In aviation, the mitigation is redundancy, because there's not an option to pull over.

Is redundancy designed into EV motors?
Wagner: Typically, redundancy is not designed into electric-vehicle drive systems solely for the purpose of redundancy. There are some cars now that have all-wheel drive—so there's a motor on the front, a motor on the back—so as a secondary feature you get the redundancy. But it wasn't done with the primary intent of having redundancy.

How does Joby's eVTOL manufacturing compare to EV manufacturing?
Wagner: The most efficient way to run a large-scale engineering effort in a mature industry, such as automotive, is to break your system up into pieces that can be outsourced to suppliers who are going to do a really good job on each piece. The downside is that when you break a problem up into three pieces, you now have interface boundaries between each of these pieces, and those always create inefficiencies. We were able to design highly integrated solutions without taking that manufacturing penalty.

Are there any materials you're really excited about?
Wagner: Permendur [a cobalt-iron alloy] typically costs about 10 times as much as traditional motor steel. That's significant, and it's often not used in ground transportation because of that cost. It comes with small improvements in performance, but enough that for aviation it's quite interesting.

Will electric aircraft catch on like ground EVs?
Wagner: I've always wanted to be a very forward thinker with respect to power train. However, one of the things I've learned over the years is that power-train development has to come with a very healthy dose of patience. Developing a whole new type of power train is a big endeavor, but it's one that I'm very confident the aviation industry will undertake. We're certainly undertaking it here at Joby, and we'll see that broaden, I'm sure, with time. ■

JOBY AVIATION

CYBORG TECH

From the Inside

CLOCKWISE FROM TOP RIGHT: 60 MINUTES/CBS NEWS; ANDREW SPEAR/REDUX; RODNEY DECKER; IAN BATES/THE NEW YORK TIMES/REDUX; BRYAN ANSELM/REDUX; STEVE CRAFT/GUARDIAN/EYEVINE/REDUX

TECHNOLOGIES THAT MERGE with the human body are often judged by the discrete tasks they can perform: restore movement, enable communication, control robotic tools. But their true test is in how they work in daily life, on and within the people who use them. In this special report, we look at exoskeletons and brain-computer interfaces from the user's perspective, where performance is measured not just in capability but in effort, trade-offs, and lived experience. → → →

Robert Woo has spent 15 years testing and improving exoskeletons from the inside

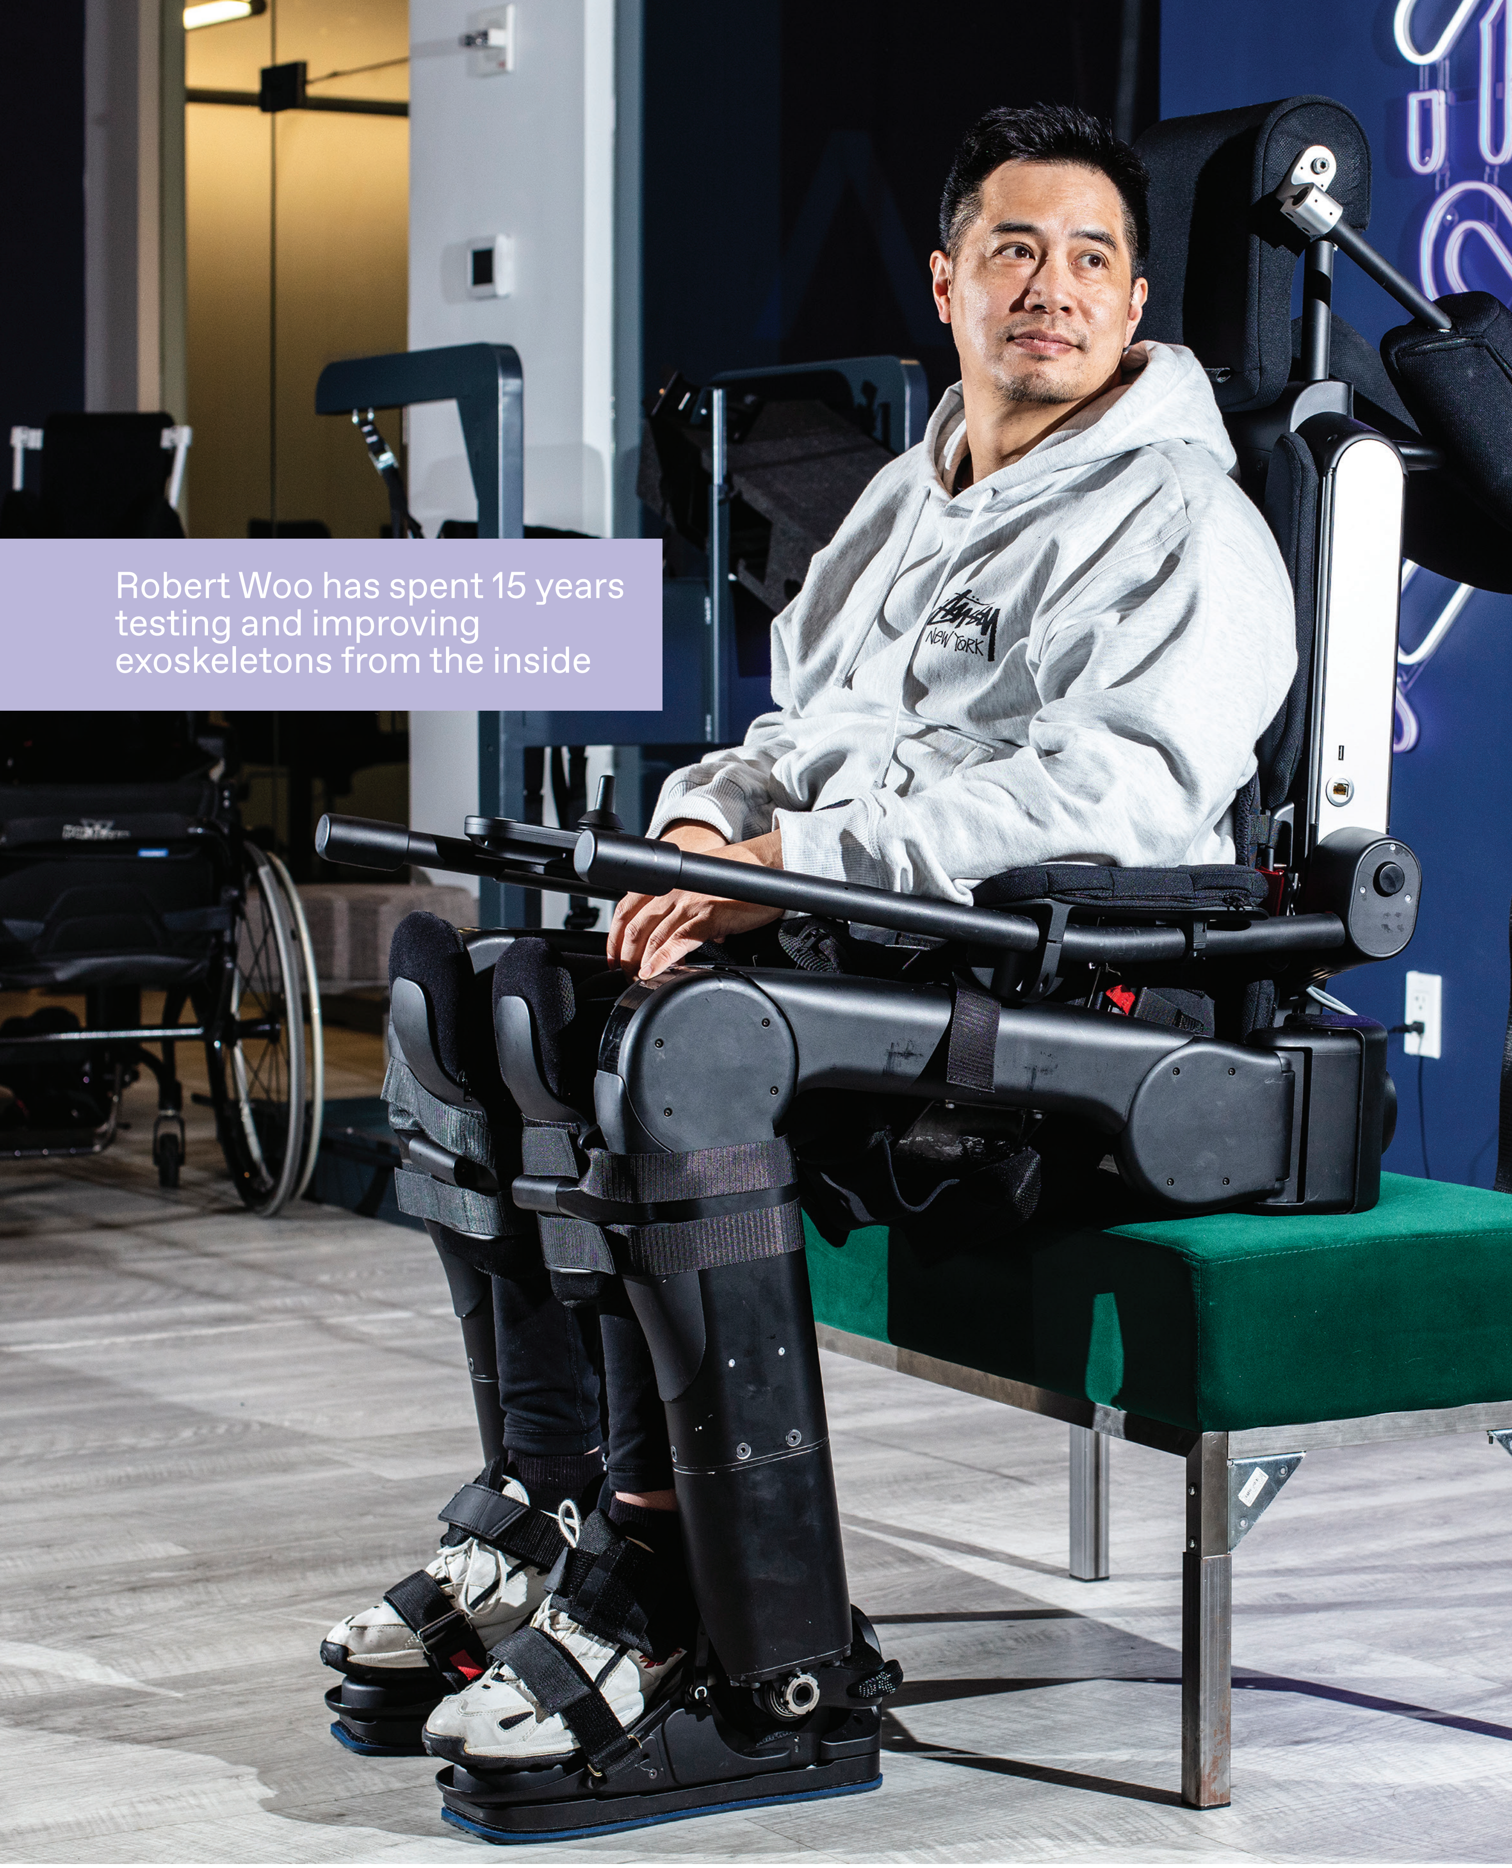

Woo prepares to test out a self-balancing exoskeleton from Wandercraft, one of the first systems designed to let users walk without crutches.

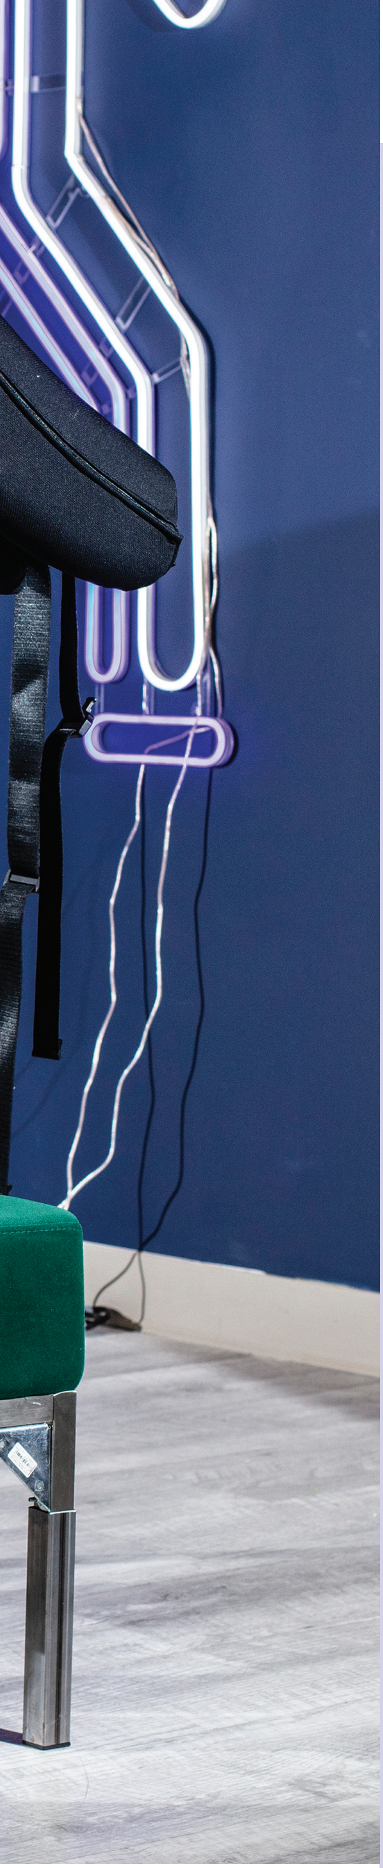

Photo by Bryan Anselm/Redux

Cyborg Tech From the Inside:

EXOSKELETONS

By **ELIZA STRICKLAND**

The MAN in the Machine

IT'S EASY TO ASSUME that Robert Woo was defined by the accident that took away his ability to walk. ¶ Certainly, the day of his accident—14 December 2007—was a turning point. Woo, an architect working on the new Goldman Sachs headquarters in New York City, hadn't attended his company's holiday party the night before, and that morning he was the only one in the trailer that served as the construction-site office. He was bent over his laptop when, 30 floors above, a crane's nylon sling gave way, sending about 6 tonnes of steel plummeting toward the trailer. The roof collapsed, folding Woo in half and smashing his face into his laptop, which smashed through his desk. ¶ "I was conscious throughout the whole ordeal," Woo remembers. "It was an out-of-body experience. I could hear myself screaming in pain. I could hear the voices of the rescue workers. I heard one firefighter say, 'Don't worry, we're getting to you.'" The rescue workers hauled him out of the rubble and got him to the emergency room in 18 minutes flat; with one lung crushed and the other punctured, he wouldn't have lasted much longer. In those frantic early moments, a doctor told him that he might be paralyzed from the neck down for the rest of his life. He remembers asking the doctors to let him die.

Woo simply couldn't imagine how a paralyzed version of himself could continue living his life. Then 39 years old, he worked long hours and jetted around the world to supervise the construction of skyscrapers. More important, he had two young boys, ages 6 months and 2 years. "I couldn't see having a life while being paralyzed from the neck down, not being able to teach my boys how to play ball," he recalls. "What kind of life would that be?"

But in a Manhattan showroom last May, Woo showed that he's not defined by that accident, which left him paralyzed from the chest down, but with the use of his arms. Instead, he has defined himself by how he has responded to his injury, and the new life he built after it.

In the showroom, Woo transferred himself from his wheelchair to an 80-kilogram (176-pound) exoskeleton suit. After strapping himself in, he manipulated a joystick in his left hand to rise from a chair and then proceeded to walk across the room on robotic legs. Woo's steps were short yet smooth, and he clanked as he walked.

This exoskeleton, from the French company Wandercraft, is one of the first to let the user walk without arm braces or crutches, which most other models require to stabilize the user's upper body. The battery-powered exoskeleton took care of both propulsion and balance; Woo just had to steer. The bulky apparatus had a backplate that extended above Woo's head, a large padded collar, armrests, motorized legs, and footplates. Walking across the room, he appeared to be half man, half machine. On the other side of the showroom's plate-glass window, on Park Avenue, a kid walking by with his family came to a dead halt on the sidewalk, staring with awe at the cyborg inside.

The amazement on the boy's face was reminiscent of Woo's young sons' reaction when they saw a photo of Woo trying out an early exoskeleton, back in 2011. "Their first comment was, 'Oh, Daddy's in an Iron Man suit,'" he remembers. Then they asked, "When are you going to start flying?" To which Woo replied, "Well, I've got to learn how to walk first."

The title of exoskeleton superhero suits Woo. He's as soft-spoken and mild-mannered as Clark Kent, with a smile that lights up his face. Yet the strength underneath is undeniable; he has built a new life out of sheer determination. For 15 years, he's been a test pilot, early adopter, and clinical-study subject for the most prominent exoskeletons under development around the world. He

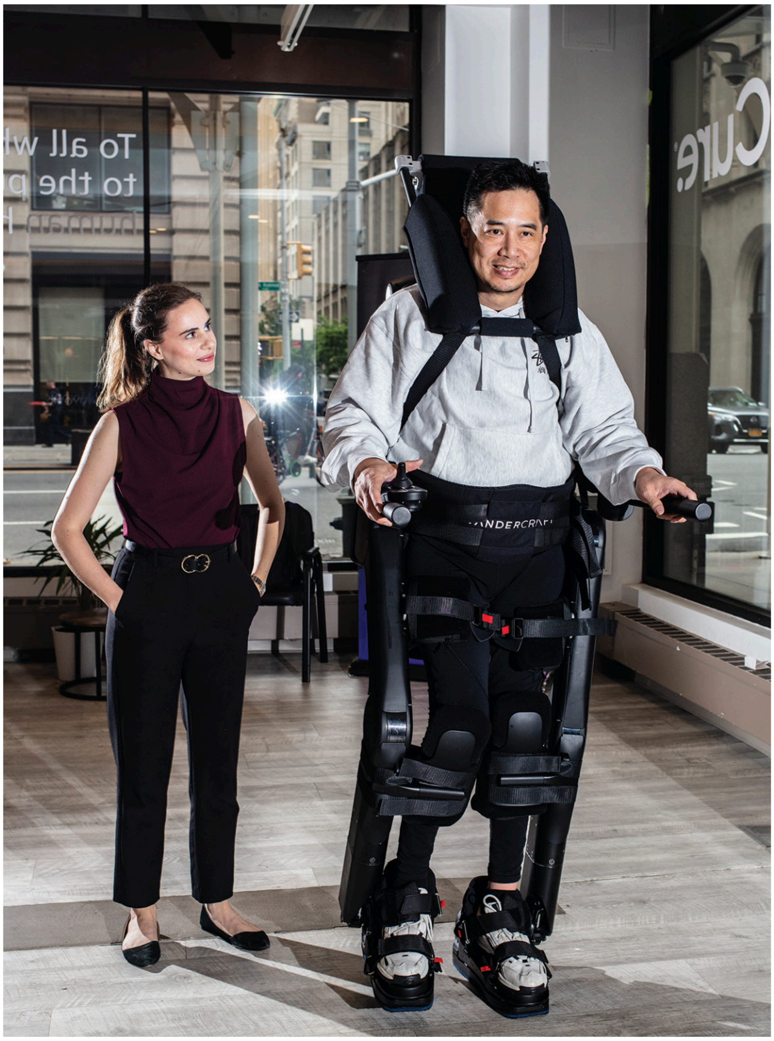

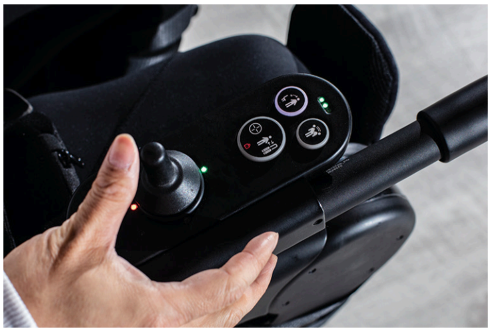

Woo walks with the Wandercraft exoskeleton as a physical therapist stands by. With the device's controller, he can stand up, initiate walking, and choose a direction.

Photos by Bryan Anselm/Redux

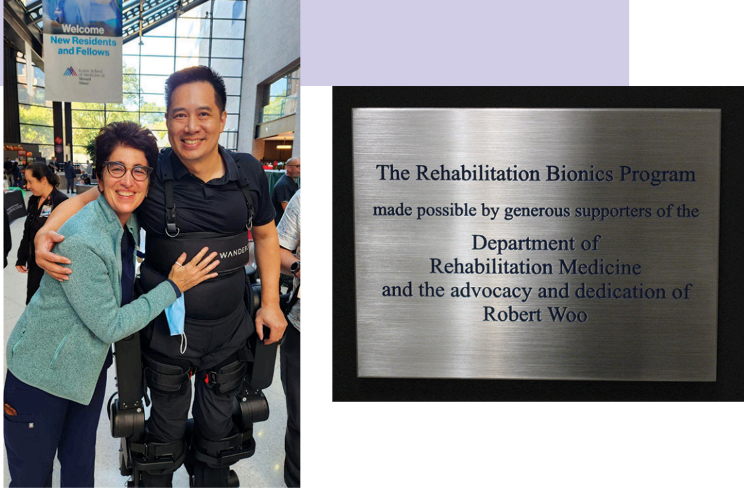

The Rehabilitation Bionics Program
made possible by generous supporters of the
Department of
Rehabilitation Medicine
and the advocacy and dedication of
Robert Woo

Angela Riccobono of Mount Sinai Hospital credits Woo with launching the hospital's bionics program; a plaque in the hospital recognizes his role.

placed the first order for an exoskeleton that was approved for home use, and he learned what it was like to be Iron Man around the house. Throughout it all, he has given the companies detailed feedback drawn from both his architectural design skills and his user experience. He has shaped the technology from inside of it.

Saikat Pal, a researcher at the New Jersey Institute of Technology, in Newark, met Woo during clinical trials for Wandercraft's first model. Like so many others in the field, Pal quickly recognized that Woo brought a lot to the table. "He's a super-mega user of exoskeletons: very enthusiastic, very athletic," Pal says. "He's the perfect subject."

By pushing the technology forward, Woo has paved the way for thousands of people with spinal cord injuries as well as other forms of paralysis, who are now benefiting from exoskeletons in rehab clinics and in their homes. "Our bionics program at Mount Sinai started with Robert Woo," says Angela Riccobono, the director of rehabilitation neuropsychology at Mount Sinai Hospital, in New York City, where Woo became an outpatient after his accident. "We have a plaque that dedicates our bionics program to him."

It's a fitting tribute. Woo's post-accident life has been marked by victories, frustrations, deep love, and one devastating loss, and yet he has continued to devote himself to bionics. And while his vision for exoskeletons hasn't changed, experience has reshaped what he expects from them in his lifetime.

Long before Woo ever stood up in a robotic suit, he had developed the habits of mind that would later make him an unusually perceptive test pilot.

Woo has always been a builder, a tinkerer, a fixer. Growing up in the suburbs of Toronto, he put together model kits of battleships and airplanes without looking at the instructions. "I just put things together the way I thought it would work out," he says. He trained as an architect and in 2000 joined the Toronto-based firm Adamson Associates Architects, a job that soon had him traveling to Europe and Asia to work on corporate high-rises.

Adamson specializes in taking the stunning designs of visionary architects and turning them into practical buildings with elevators and bathrooms. "Most of the design architects don't really have a clue about how to build buildings," Woo says. He liked solving those problems; he liked reconciling beautiful designs with the stubborn reality of construction. That talent for understanding a structure from the inside and spotting the flaws would prove essential later.

After his accident, Woo had two major surgeries to stabilize his crushed spine, which required surgeons to cut through muscles and nerves that connected to his arms. For two months, he couldn't feel or move his arms; there was a chance he never would again. Only when sensation began creeping back into his fingertips did he allow himself to imagine a different future. If he wasn't paralyzed from the neck down, he thought, maybe more of his body could be brought back on line. "My focus was to walk again," he says.

Woo was discharged in March 2008 and went back to his New York City apartment. He was still bedridden and required around-the-clock care. He doesn't much like to talk about this next part: By May, his then-wife had moved back to Canada and filed for divorce, asking for full custody of their two children. Woo remembers her saying, "I can't look after three babies, and one of them for life."

It was a dark time. Riccobono of Mount Sinai, who met Woo shortly after he became an outpatient there in 2008, recalls the despondent look on his face the first time they talked. "I wasn't sure that he wasn't going to take his life, to be honest," she says. "He felt like he had nothing to live for."

Yet Woo harbors no animosity toward his ex-wife. "If we hadn't separated and gone through the custody hearing, I don't think I would have gotten this far," he says. To win partial custody of his children, Woo had to become independent. He had to get off narcotic pain medications, regain strength, and learn how to navigate life in a wheelchair. He had to show that he no longer needed constant nursing, and that he could take care of both himself and his boys.

There were milestones: learning how to get back into his wheelchair after a fall, learning to drive a car with hand controls, learning to manage his body as it was, not as it had been. The biggest change came when he reconnected with his high school sweetheart, a vivacious woman named Vivian Springer. She was then dividing her time between Toronto and New York City, and she had a son who was almost the same age as Woo's two boys. Springer had worked in a nursing home and knew how to change the sheets without getting him out of bed; she was currently working in human resources and knew how to deal with insurance companies. "You wouldn't believe how much stress it lifted off of me," Woo says. Over time, they became a family.

Once Woo had that foundation in place, Riccobono witnessed a profound change. "He went from focusing on 'What I can't do anymore' to 'What's still possible? What can I do with what I have?'" At Mount Sinai, Woo remembers asking his doctor Kristjan Ragnarsson, who was then chairman of the department of rehabilitation medicine, if he

ROBERT WOO

would ever walk again. "His response was, 'Yes, you can walk again,'" Woo remembers, "'but not the way you used to walk.'"

As soon as he had regained use of his hands, Woo had started googling, looking for anything that could get him back on his feet. He tried rehab equipment like the Lokomat, which uses a harness suspended above a treadmill to enable users to walk. But at the time, it required three physical therapists: one to move each leg and one to control the machine. It was a far cry from the independent strides he dreamed of.

Several years in, he learned about two companies that had built something radically different: exoskeleton suits for people with spinal cord injuries. These prototypes had motors at the knees and the hips to move the legs, with the user stabilizing their upper body with arm braces. Woo desperately wanted to try one, although the technology was still experimental and far from regulatory approval. So he took the idea to Ragnarsson, asking if Mount Sinai could bring an exoskeleton into its rehab clinic for a test drive. Ragnarsson, who's now retired, remembers the request well. "He certainly gave us the kick in the behind to get going with the technology," he says.

Ragnarsson had seen decades of failed attempts to get paraplegics upright, including "inflatable garments made of the same material the astronauts used when they went to the moon," he says. All those devices had proved too tiring for the user; in contrast, the battery-powered exoskeletons promised to do most of the work. And he knew the CEO of Ekso Bionics, a Berkeley, Calif.–based company that had built exoskeletons for the military. In 2011, Ekso brought its new clinical prototype to Mount Sinai.

The day came for Woo's first walk. "I was excited, and I was also scared, because I hadn't stood up for almost five years," he remembers. "Standing up for the first time was like floating, because I couldn't feel my feet." In that first Ekso model, Woo didn't control when he stepped forward; instead, he shifted his weight in preparation, and then a physical therapist used a remote control to trigger the step. Woo walked slowly across the room, using a walker to stabilize his upper body, his steps a symphony of clunks and creaks and whirs. He found it mentally and physically exhausting, but the effort felt like progress.

At Mount Sinai Hospital, Woo tries an Ekso Bionics device as his doctor, Kristjan Ragnarsson [right], looks on.

MARIO TAMA/GETTY IMAGES

Riccobono was there for those first steps, with tears running down her face. "I remembered how he looked the day I first met him, so defeated," she says. "To see him rise from the chair, to see him rise to a standing position, to see how tall he was, to see him take those first steps—it was beautiful." Ragnarsson saw clear benefits to the technology. "Any type of walking is good physiologically," he says. "And it's a tremendous boost psychologically to stand up and look someone in the eye." Woo remembers hugging his partner, Springer, and for the first time not worrying about running over her toes with his wheelchair. I first met Woo a few days later, during his third session with the Ekso at Mount Sinai.

Later that same year, at a Department of Veterans Affairs (VA) hospital in the Bronx, Woo got to try a prototype of the world's other leading exoskeleton: the ReWalk, from the Israeli company of the same name (since renamed Lifeward). VA researchers, led by Ann Spungen, were keen to determine if exoskeleton use had real medical value for people with spinal cord injuries. Woo was part of that clinical trial, for which he had more than 70 walking sessions, and he's since been in many others. But he remembers the first VA trial with the most gratitude. "Dr. Spungen's first exoskeleton clinical trial really turned things around for me," he says.

Over the course of the trial's nine intense months, Woo says he saw noticeable improvements to many facets of his health. "By the end of the trial, I eliminated about three-quarters of my medication intake," he says, including narcotic pain pills and medication for muscle spasms. He grew fitter, with less body fat, more muscle mass, and lower cholesterol. His circulation improved, he says, causing scrapes and cuts to heal more quickly, and his digestion improved too. The results Woo experienced have generally been borne out in research studies at the VA and elsewhere—exoskeletons aren't just good for the mind, they're good for the body.

During the VA trial, Woo began to think of exoskeletons not as miraculous machines but as works in progress.

Pierre Asselin, the biomedical engineer coordinating the VA's study, watched participants respond very differently to the equipment. "These devices are not the equivalent of walking—

Woo's sons, Tristan [far left] and Adrien, grew up watching their dad test exoskeletons. His wife, Vivian, was trained in how to operate the device he used at home.

LEFT: ROBERT WOO; RIGHT: LIFEWARD

you're tired after walking a mile," he says. He notes that later models of both the Ekso and ReWalk enabled users to initiate each step through software that recognized when they shifted their weight. Asselin adds that the cognitive load is "like learning to drive a manual transmission car, where at first you're really struggling to coordinate the clutch and the brake." Woo picked it up immediately, he remembers.

Woo became an invaluable partner, Asselin says. "When we first started with the devices, there was no training manual. We developed all of that through collaboration with Robert and other participants." Woo pushed the limits of the technology, Asselin says, whether it was seeing how many steps he could take on one battery charge or simulating a failure mode. "He'd say, 'What happens if I was to fall? What would be the approach to getting up?'"

Woo approached the ReWalk the way he had approached buildings in his previous life: He looked inside the structure and found the weak points. An early model left some users with leg abrasions where the straps rubbed—a small injury for most people but a serious risk for someone who can't feel a wound forming. Woo suggested better padding and stronger abdominal supports to redistribute the load. He also hated the heavy backpack that carried the battery and computer, so one afternoon he grabbed an old pack, cut off the straps, and rebuilt it into a compact hip-mounted pouch. Then he snapped photos and sent them to the company. The next model arrived with a fanny pack.

Sometimes his fixes were more ambitious. One Ekso unit that he used at Mount Sinai kept shutting down after 30 minutes. Woo felt the hip motors and found them hot to the touch. "I said, 'Can I remove these? I'm going to make a really quick fix, okay? Give me a drill and I'll put a couple of holes in it," he recalls telling the therapists, proposing to create a DIY heat sink. He wasn't allowed to modify the prototype, but a year later the company introduced improved cooling around the hip motors. "There is a Robert Woo design on this device," one therapist told him.

Eythor Bender, who was then the CEO of Ekso, called Woo to thank him for his feedback and invite him to spend a week at Ekso's headquarters. "There was no lack of engineering power in that building," says Bender. "But sometimes when you work with engineers, they overlook important things." Bender says Woo brought both design skills and lived experience to his weeklong residency. "He told the engineers, 'Guys, this has to be something that people actually like to wear.'"

The longer Woo tested, the further ahead he started thinking. With motors only at the hips and knees, every exoskeleton still required crutches. Add powered ankles, he told the Ekso and ReWalk teams, and the suits could balance themselves, freeing the user's hands. But Woo was ahead of his time. "They said they weren't going to do that. They weren't going to change their whole platform," he remembers. Years later, though, hands-free exoskeletons like those from Wandercraft would emerge built around exactly that principle.

By the mid-2010s, Woo had pushed the technology as far as he could in clinics. What he wanted now was to use an exoskeleton at home.

That milestone came after ReWalk's exoskeleton became the first to win FDA approval for home use in 2014. ReWalk engineers still remember Woo's help on the final tests for that personal-use model. It was the end of May in 2015, recalls David Hexner, the company's vice president of research and development. "He said, 'Guys, this is great. I'm going to buy it.'"

Woo was the first customer to buy an exoskeleton to bring home, paying US $80,000 out of pocket. His insurance wouldn't cover the cost, but he was able to make the purchase in part because of a legal settlement after his accident. The home-use model came with a requirement that the user have at least one companion who was fully trained in operating the device. In Woo's case, that meant that Springer learned to suit him up, realign his balance, and help him if he fell.

On delivery day, two SUVs drove up to a hotel down the street from Woo's condo in the Toronto area. The technicians hauled two huge boxes into a hotel room and assembled his personal exoskeleton. They took Woo's measurements, made adjustments, checked the software. This latest version could be controlled by either weight shifting or tapping commands on a smartwatch, and Woo had the app ready. He tested out everything in the hotel room, signed off, and then the technicians drove his robot legs to his home.

That was the start of his golden period with the ReWalk—similar to the excitement many people experience with a new piece of exercise equipment. "I used it every day for a few hours, and then I started logging how many steps I'd done," Woo says. "My last count was probably just slightly over a million steps," he says, with half of those steps taken in his home unit and half in training programs and clinical trials.

Tristan, Woo's eldest son, remembers doing laps with his dad in the condo's underground parking garage while his dad was training for a 5-kilometer race in New York City. Tristan admits that he had previously been embarrassed about his dad, but training for the race shifted something for him. "I was so used to not wanting to tell people that my dad was in a wheelchair, but then I shared his passion for the training," he says. "When people would come up to us, I'd tell them about it."

The ReWalk could turn ordinary moments into small engineering projects. On weekends, Woo would take his boys to the golf course behind their condo and bring a baseball. He had rigged two holsters to the sides of the suit so he could stash a crutch and stand on three points (two legs and one arm) while he pitched or caught. Throw, switch crutches, catch. On the day of his accident, he never thought such a scene would be possible. But with the exoskeleton, it became just another design problem to solve. "It's a little more work. It's not perfect," he says. "But in the end, you still get to do what you want to do—which is play ball with your sons."

Tristan, now a college student, says he didn't realize at the time how hard his dad worked to make those mundane activities possible. "Reflecting on it now," he says, "he has shaped almost every element of my life, and he definitely is my hero."

But even during that golden stretch, the ReWalk had a way of asserting its limits. Every so often it would freeze mid-stride and require a reboot—a small technical hiccup in theory, but a serious problem when there's a person strapped inside. Once, when he was walking on his own in the parking garage (without his mandated companion), the suit glitched and went into "graceful collapse" mode, lowering him to a seated position on the ground. Woo had to ask security to bring his wheelchair and a dolly.

He had imagined the exoskeleton would be most useful in the kitchen. Woo loves to cook, and he had pictured himself standing at the stove, looking down into pots, and moving easily between counter and sink. The reality, he found out, was more complicated. "It's actually very time-consuming and troublesome" to cook in an exoskeleton, he says.

Preparing a meal meant first rolling through the kitchen in his wheelchair to gather every ingredient and utensil, then transferring himself into the ReWalk and moving himself into position at the counter, stopping at just the right moment. "That's when I fell once," Woo says. "I collided with the counter and then lost my balance and fell backward." If all went well, he'd lean either on one crutch or the counter to keep his balance while he worked. But if he'd forgotten to grab the vinegar from the cabinet, he'd have to go into walk mode, crutch over to it, and figure out how to carry the bottle back to his workstation.

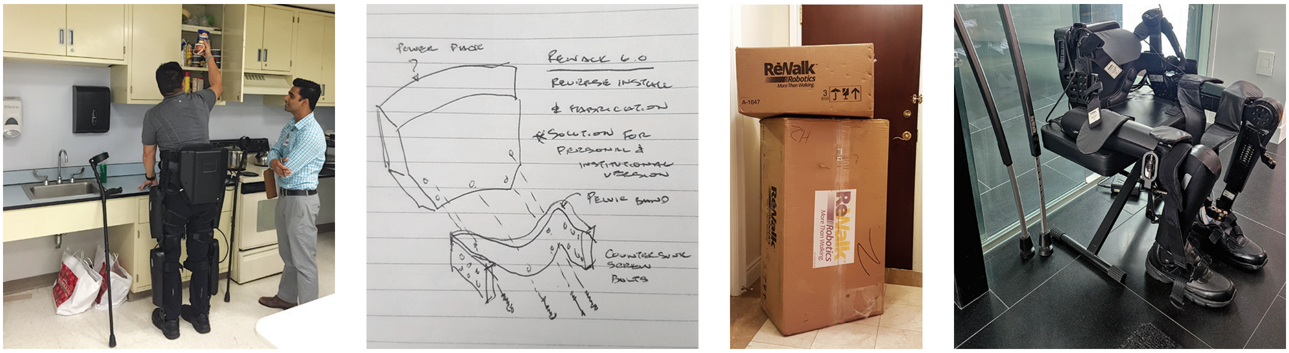

At Mount Sinai Hospital in 2015, Woo tried using an exoskeleton in a kitchen [left]. He often sent detailed diagrams to the companies with design suggestions [middle left]. His personal ReWalk arrived in two big boxes [middle right], but today it's collecting dust on a bench in his house [right].

FROM LEFT: ELIZA STRICKLAND; ROBERT WOO (3)

Gradually, he stopped trying. The suit, which he'd once worn every day, spent more time sitting idle in the hallway; like so many abandoned treadmills and stationary bikes, it gathered dust. Part of the reason was the exoskeleton's practical limitations, but part of it was a shocking development: In 2024, Vivian was diagnosed with an aggressive form of breast cancer. She died in November of that year, at the age of 54.

Woo was scheduled to begin a new round of clinical trials for the Wandercraft home-use exoskeleton that month. In the aftermath of Vivian's death, he postponed his sessions and questioned whether he would ever go back. "At the time, I thought, 'What's the point?'" he remembers.

He did go back, though. "He just rolled up, right into my office," says Mount Sinai's Riccobono. "He still had Vivian's box of ashes on his lap. That's how fresh it was." Woo brought the box into a meeting of spinal cord injury patients and shared the story of losing the love of his life. And he told them that he heard his wife's voice in his head every day, telling him to get back to work. Once again, he was figuring out how to move forward with what he had.

In the Wandercraft showroom last May, Woo steered toward the door to the street, technicians flanking him like spotters. The slope down to the sidewalk was barely an inch high, but everyone tensed. He shifted his weight and took a step forward. The suit halted automatically. He tried again—step, stop; step, stop—as the suit kept detecting the slight decline and a safety feature kicked in. The Wandercraft isn't yet rated for slopes of more than 2 percent, and even the gentle pitch of Park Avenue was enough to trigger its safeguards. When he finally reached the sidewalk, Woo broke into a grin. A man in the back seat of a stopped Uber leaned out his window, filming.

Woo had recently completed seven sessions with the Wandercraft at the VA hospital and had been impressed overall. But at the showroom, he rolled up his pants leg to reveal an abrasion on his shin, the result of a strap that had worn away a patch of skin during a long walking session. He would later send Wandercraft a nine-page assessment with photos and a technology wish list, asking the company to work on things like padding, variable walking speeds, and deeper squats.

Wandercraft's engineers relish that kind of user feedback, says CEO Matthieu Masselin. Exoskeletons are a far more difficult engineering problem than humanoid robots, he explains. "You basically have two systems of equal importance. You know about the robot—it's fully quantified and measured. But you don't know what the person is doing, and how the person is moving within the device."

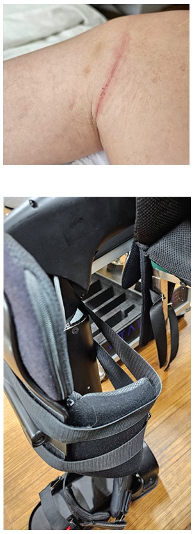

During testing of the Wandercraft exoskeleton, straps caused an abrasion on Woo's leg, an issue he documented in feedback to the company.

ROBERT WOO

Since Woo began testing exoskeletons 15 years ago, both the technology and the market have made strides. ReWalk and Ekso won FDA clearance for clinical use in the 2010s, and both now sell home-use versions. The companies have sold thousands of exoskeletons to rehab clinics and personal users, and they see room for growth; in the United States alone, about 300,000 people live with spinal cord injuries, and millions more have mobility impairments from stroke, multiple sclerosis, or other conditions. The VA began supplying devices to eligible veterans in 2015, and Medicare recently established a system for reimbursement, a move that private insurers are beginning to follow. What was once experimental is slowly becoming established.

Researchers who test the devices say the technology still has significant limits. Pal, of the New Jersey Institute of Technology, mentions battery life, dexterity, and reliability as ongoing challenges. But, he says with a laugh, "Our bodies have evolved over many millions of years—these machines will need a bit more time." Pal hopes the companies will keep pushing the technological frontier. "My lifetime goal is to see the day when someone like Robert Woo can wake up in the morning, put this device on, and then live an ordinary life."

For Woo, the real question about the self-balancing Wandercraft was: Could he cook with it? In the VA hospital's home mock-up, he tried it out in the kitchen, stepping sideways to retrieve items from cabinets and squatting to grab something from the fridge's lower shelf. For the first time in years, he could work at a counter without leaning on crutches. "The self-standing exoskeleton changes everything," he says. He imagines a user placing a Thanksgiving turkey on a tray attached to the suit and walking it into the dining room.

Back in the showroom, Woo finishes the demo and brings the suit to a seated position before transferring back to his wheelchair. After so many years of testing prototypes, he's now realistic about the technology's timeline. A truly all-day exoskeleton—the kind you live in, the kind that replaces a wheelchair—may be a decade or more away. "It may not be for me," he says. But that's no longer the point. He's thinking about young people who are newly injured, who are lying in hospital beds and trying to imagine how their lives can continue. "This will give them hope." ■

HOW THIS STORY WAS REPORTED The author has followed Robert Woo's use of exoskeletons since 2011, observing early Ekso trials at Mount Sinai, a ReWalk demo on the streets of Manhattan in 2015, further sessions at the VA and Mount Sinai in 2015, and 2025 testing of the Wandercraft.

Cyborg Tech From the Inside:

BRAIN IMPLANTS

By **EDD GENT**

Life With an Experimental Brain Implant

Early users reveal what the technology gives—and what it takes

SCOTT IMBRIE VIVIDLY REMEMBERS the first time he used a robotic arm to shake someone's hand and felt the robotic limb as if it were his own. "I still get goosebumps when I think about that initial contact," he says. "It's just unexplainable." The moment came courtesy of a brain implant: an array of electrodes that let him control a robotic arm and receive tactile sensations back to the brain. ¶ Getting there took decades. In 1985, Imbrie had woken up in the hospital after a car accident with a broken neck and a doctor telling him he'd never use his hands or legs again. His response was an expletive, he says—and a decision. "I'm not going to allow someone to tell me what I can and can't do." With the determination of a head-strong 22-year-old, Imbrie gradually regained the ability to walk

STEVE CRAFT/GUARDIAN/EYEVINE/REDUX

NOLAND ARBAUGH

Arbaugh, the first recipient of Neuralink's brain-computer interface, says that using the implant to control a computer made him feel independent and free.

and some limited arm movement. Aware of how unusual his recovery was, the Illinois native wanted to help others in similar situations and began looking for research projects related to spinal cord injuries. For decades, though, he wasn't the right fit, until in 2020 he was finally accepted into a University of Chicago trial.

Imbrie is part of a rarefied group: More people have gone to space than have received advanced brain-computer interfaces (BCI) like his. But a growing number of companies are now attempting to move the devices out of neuroscience labs and into mainstream medical care, where they could help millions of people with paralysis and other neurological conditions. Some companies even hope that BCIs will eventually become a consumer technology.

None of that will be possible without people like Imbrie. He's a member of the BCI Pioneers Coalition, an advocacy group founded in 2018 by Ian Burkhart, the first quadriplegic to regain hand movement using a brain implant.

That life-changing experience convinced Burkhart that BCIs will make the leap from lab to real world only if users help shape the technology by sharing their perspectives on what works, what doesn't, and how the devices fit into daily life. The coalition aims to ensure that companies, clinicians, and regulators hear directly from trial participants.

The group also serves as a peer-support network for trial participants. That's crucial, because despite the steady drumbeat of miraculous results from BCI trials, receiving a brain implant comes with significant risks. Surgical complications, such as bleeding or infection in the brain, are possible. Even more concerning is the potential psychological toll if the implant fails to work as expected or if life-changing improvements are eventually withdrawn.

Researchers spell this out up front, and many are put off, says John Downey, an assistant professor of neurological surgery at the University of Chicago and the lead on Imbrie's clinical trial. "I would say, the number of people I talk to about doing it is probably 10 to 20 times the number of people that actually end up doing it," he says.

BCI PIONEERS ARRIVE at their unique status via a number of paths, including spinal cord injuries, stroke-induced paralysis, and amyotrophic lateral sclerosis (ALS). The implants they receive come from Blackrock Neurotech, Neuralink, Synchron, and other companies, and are being tested for restoring limb function, controlling computers and robotic arms, and even restoring speech.

Many of the implants record signals from the motor cortex—the part of the brain that controls voluntary movements—to move external devices. Some others target the somatosensory cortex, which processes sensory signals from the body, including touch, pain, temperature, and limb position, to re-create tactile sensation.

Ease of use depends heavily on the application. Restoring function to a user's own limbs or controlling robotic arms involves the most difficult learning curve. In early sessions, participants watch a virtual arm reach for objects while they imagine or attempt the same movement. Researchers record related brain signals and use them to train "decoder" software, which translates neural activity into control signals for a robotic arm or stimulation patterns for the user's nerves or muscles.

Paralyzed in a 2010 swimming accident, Burkhart took part in a trial conducted by Battelle Memorial Institute and Ohio State University from 2014 to 2021. His implant recorded signals from his motor cortex as he attempted to move his hand, and the system relayed those commands to electrodes in his arm that stimulated the muscles controlling his fingers.

Getting the system to work seamlessly took time, says Burkhart, and initially required intense concentration. Eventually, he could shift his focus from each individual finger movement to the overall task, allowing him to swipe a credit card, pour from a bottle, and even play *Guitar Hero*.

Training a decoder is also not a one-and-done process. Systems must be regularly recalibrated to account for "neural drift"—the gradual shift in a person's neural activity patterns over time. For complex tasks like robotic arm control, researchers may have to essentially train an entirely new decoder before each session, which can take up to an hour.

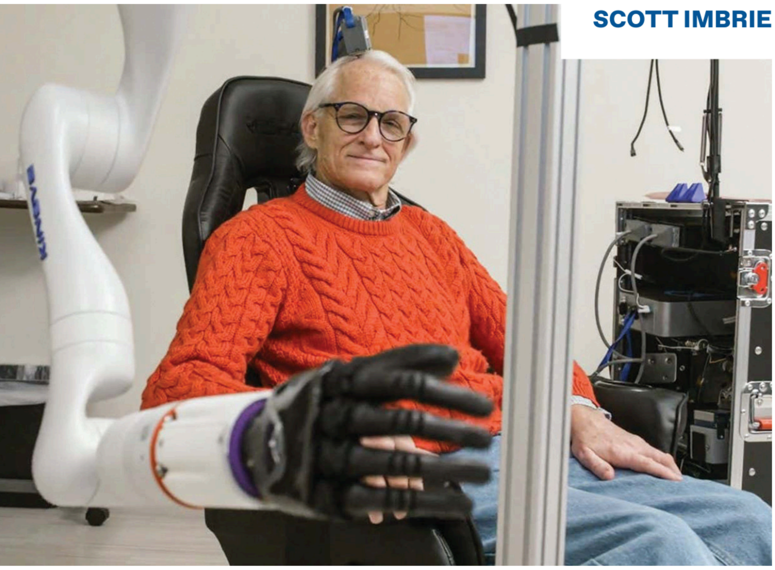

SCOTT IMBRIE

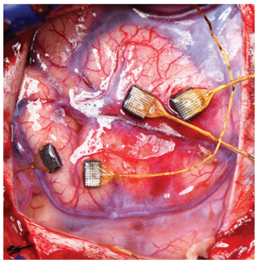

Imbrie has shaken hands with a robotic arm controlled by a brain implant. The electrodes record neural signals that enable him to move the device and receive tactile feedback.

LEFT: 60 MINUTES/CBS NEWS; RIGHT: UNIVERSITY OF CHICAGO

IAN BURKHART

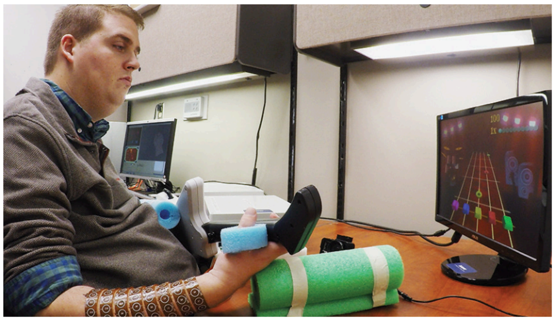

Burkhart, who is paralyzed from the chest down, received a brain implant that routed neural signals through a computer to his paralyzed muscles, enabling him to grasp objects and even play the video game *Guitar Hero*.

CLOCKWISE FROM LEFT: ANDREW SPEAR/REDUX; IAN BURKHART (2); BATTELLE

Even after the system is ready, using the device can be taxing, says Austin Beggin, who was paralyzed in a swimming accident in 2015 and now participates in a Case Western Reserve University trial aimed at restoring hand movement. “The mental work of just trying to do something like shaking hands or feeding yourself is 100-fold versus you guys that don’t even think about it,” he says.

It’s also a serious time commitment. Beggin travels over 2 hours from his home in Lima, Ohio, to Cleveland for two weeks every month to take part in experiments. All the equipment is set up in the house he stays in, and he typically works with the researchers for 3 to 4 hours a day. The majority of the experiments are not actually task-focused, he says, and instead are aimed at adjusting the control software or better understanding his neural responses to different stimuli.

But the BCI users say the hard work is worth it. Beyond the hope of restoring lost function, many feel a strong moral obligation to advance a technology that could help others. Beggin compares the pioneers to the early astronauts who laid the groundwork for the lunar landings. “We’re some of the first astronauts just to get shot up for a couple of hours and come back down to earth,” he says.

SPEAK TO BCI EARLY ADOPTERS and a pattern emerges: The biggest benefits are often more emotional than practical. Using a robotic arm to feed yourself or control a computer is clearly useful, but many pioneers say the most meaningful moments are the ones the experiment wasn’t even trying to produce. Beggin counts shaking his parents’ hands for the first time since his injury and stroking his pet dachshund as among his favorite moments. “That stuff is absolutely incredible,” he says.

Neuralink participant Alex Conley, who broke his neck in a car accident in 2021, uses his implant to control both a robotic arm and computers, enabling him to open doors, feed himself, and handle a smartphone. But he says the biggest boost has come from using computer-aided design software.

A former mechanic, Conley began using the software within days of receiving his implant to design parts that could be fabricated on a 3D printer. He has designed everything from replacement parts for his uncle’s power tools to bumpers for his brother-in-law’s truck. “I was a very big problem solver before my accident. I was able to fix people’s things,” he says. “This gives me that same little burst of joy.”

The outside world often underestimates those little wins, says Nathan Copeland, who holds the record for the longest functional brain implant. After breaking his neck in a car accident in 2004, he joined a University of Pittsburgh BCI trial in 2015 and has since used the device to control both computers and a robotic arm.

NOLAND ARBAUGH

Before receiving his Neuralink implant, Arbaugh used mouth-operated devices to control a computer. He says the BCI is more reliable and enables him to do many more things on his own.

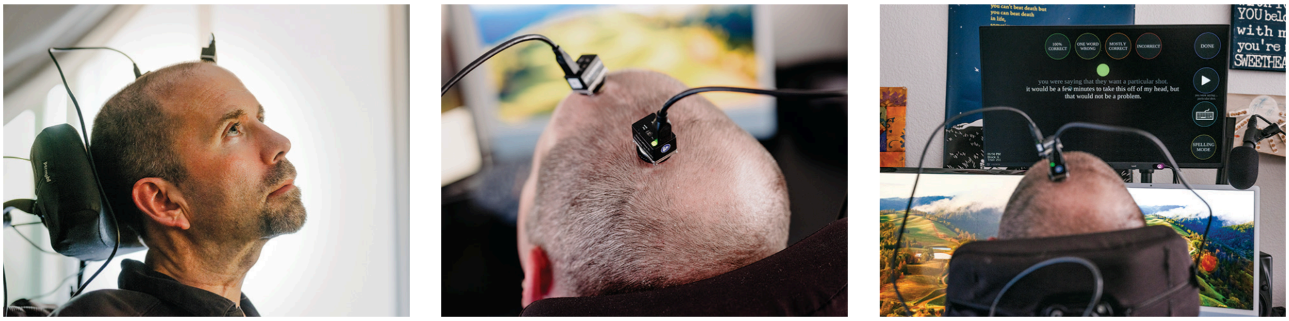

CASEY HARRELL

After being diagnosed with ALS in 2020, Harrell gradually lost the ability to speak. Surgeons implanted four Blackrock Neurotech arrays in brain regions that coordinate the muscles involved in speech.

After he uploaded a video to Reddit of himself playing *Final Fantasy XIV*, one commenter criticized him for not using his device for more practical tasks. Copeland says people don't understand that those lighthearted activities also matter. "A lot of tasks that people think are mundane or frivolous are probably the tasks that have the most impact on someone that can't do them," he says. "Agency and freedom of expression, I think, are the things that impact a person's life the most."

THIS PERSPECTIVE RESONATES with Neuralink's first user, Noland Arbaugh—paralyzed from the neck down after a swimming accident in 2016. After receiving his implant in January 2024, he was able to control a cursor within minutes of the device being switched on. A few days later, the engineers let him play the video game *Civilization VI*, and the technology's potential suddenly felt real. "I played it for 8 hours or 12 hours straight," he says. "It made me feel so independent and so free."

But the technology is also providing more practical benefits. Before his implant, Arbaugh relied on a mouth-held typing stick and a mouth-controlled joystick called a Quadstick, which uses sip-or-puff sensors to issue commands. But the fiddliness of this equipment required constant caregiver support. The Neuralink implant has dramatically increased the number of things he can do independently. He says he finds great value in not needing his family "to come in and help me 100 times a day."

For Casey Harrell, the technology has been even more transformative. Diagnosed with ALS in 2020, the climate activist had just welcomed a baby daughter and was in the midst of a major campaign, pressuring a financial firm to divest from companies that had poor environmental records.

"Every morning we'd wake up and there'd be a new thing he couldn't do, a new part of his body that didn't work," says his wife, Levana Saxon. Most alarming was his rapid loss of speech, which, among other things, left him unable to indicate when he was in pain. Then a rel-

TOP: REBECCA NOBLE/THE NEW YORK TIMES/REDUX;
BOTTOM: IAN BATES/THE NEW YORK TIMES/REDUX (3)

ative alerted him to a clinical trial at the University of California, Davis, using BCIs to restore speech. He immediately signed up.

The device, implanted in July 2023, records from the brain region that controls muscles involved in talking and translates these signals into instructions for a voice synthesizer. Within 30 minutes of it being switched on, Harrell could communicate again. “I was absolutely overwhelmed with the thought of how this would impact my life and allow me to talk to my family and friends and better interact with my daughter,” he says. “It just was so overwhelming that I began to cry.”

While earlier assistive technology limited him to short, direct commands, Harrell says the BCI is fast enough that he can hold a proper conversation, and he’s been able to resume work part-time.

BCI TECHNOLOGY still has limits. Most trial participants using Blackrock Neurotech implants can operate their devices only in the lab because the systems rely on wired connections and racks of computer hardware. Some users, including Copeland and Harrell, have had the equipment installed at home, but they still can’t leave the house with it. “That would be a big unlock if I was able to do so,” says Harrell.

The academic nature of many trials creates additional constraints. Pressure to publish and secure funding pushes researchers to demonstrate peak performance on narrow tasks rather than build more versatile and reliable systems, says Mariska Vansteensel, who runs BCI studies at the University Medical Center Utrecht in the Netherlands. She says that investigating the technology’s limits or repeating an experiment in new patients is “less rewarded in terms of funding.”

One of Imbrie’s biggest frustrations is the rapid turnover in experiments. Just as he begins to get proficient at one task, he’s asked to switch to the next task. Study designs also mean that much of the users’ time is spent on mundane tasks required to fine-tune the system.

Perhaps the biggest issue is that trials are often time-limited. That’s partly because scar tissue from the body’s immune response to the implant can gradually degrade signal quality. But constraints on funding and researcher availability can also make it impossible for users to keep using their BCIs after their trials end, even when the technology is still functional.

Burkhart has firsthand experience. His trial was extended, but the implant was eventually removed after he got an infection. He always knew the trial would end, but it was nonetheless challenging. “It was a little bit of a tease where I got to see the capability of the restoration of function,” he says. “Now I’m just back to where I was.”

PROGRESS IS BEING MADE in transitioning the technology from experimental research devices to fully-fledged medical products that could help users in their everyday lives. Most academic BCI research has relied on Blackrock Neurotech’s Utah Arrays, which typically feature 96 needle-like electrodes that penetrate the brain’s surface. The implant is connected to a skull-mounted pedestal that’s wired to external hardware. But some of the newer devices are sleeker and less invasive.

Neuralink’s implant houses its electronics and rechargeable battery in a coin-size unit connected to flexible electrode threads inserted into the brain by a robotic “sewing machine.” The implant, which is roughly

CONTINUED ON P. 50

BCI DESIGNS USED BY TODAY’S PIONEERS

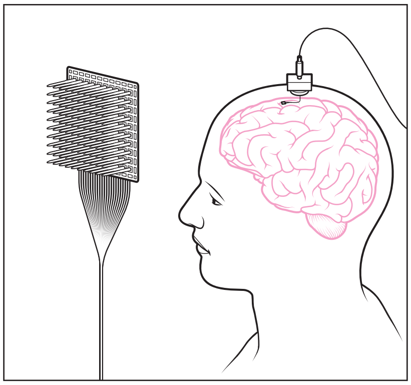

BLACKROCK NEUROTECH

The Utah Array has 96 tiny electrodes, which penetrate the brain’s surface to record neural signals. A pedestal mounted on the skull connects the system to external computers. It delivers high-quality data but often keeps users tethered to the lab.

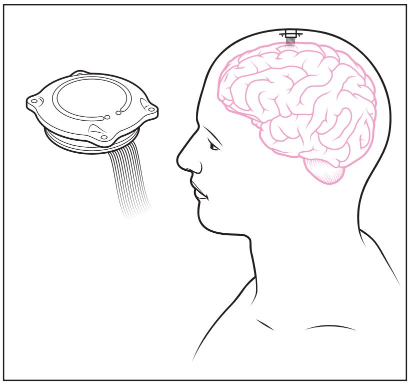

NEURALINK

Flexible electrode threads are inserted into the brain by a surgical robot; they’re connected to a coin-size implant embedded in the skull. The device transmits data wirelessly and is recharged through the skin, making it easier to use outside lab settings.

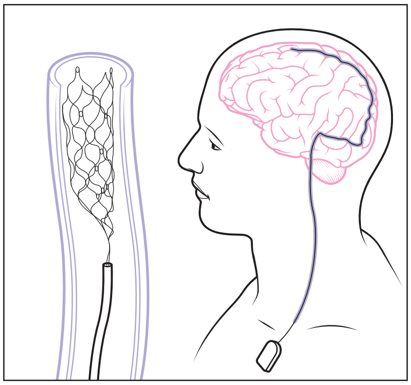

SYNCHRON

A stent-like implant is threaded through blood vessels into the brain, thus avoiding brain surgery. This Stentrode connects to a unit in the chest that powers it and transmits data wirelessly, so it can also be used outside the lab.

Illustration by Chris Philpot

Experimental Brain Implants

CONTINUED FROM P. 37

the size of a quarter or a euro, is mounted in a hole cut into the skull and charges and transfers data wirelessly. Synchron takes a different approach, threading a stent-like implant through blood vessels into the motor cortex. This Stentrode connects by wire to a unit in the chest that powers the implant and transmits data wirelessly.

Neuralink's decoder runs on a laptop, while Synchron uses a smartphone-size signal processing unit as a wireless bridge to the user's devices, which allows them to use their implants at home and on the move. The companies have also developed adaptive decoders that use machine learning to adjust to neural drift on the fly, reducing the need for recalibration.

Making these devices truly user-friendly will require technology that can interpret user context, says Kurt Haggstrom, Synchron's chief commercial officer—including mood, attention levels, and environmental factors like background noise and location. This approach will require AI that analyzes neural signals alongside other data streams such as audio and visual input.

Last year, Synchron took a first step by pairing its implant with an Apple Vision Pro headset. When Australian trial participant Rodney Gorham looked at devices such as a fan, a smart speaker, and an air conditioner, the headset overlaid a menu that let him adjust the device's settings using his implant.

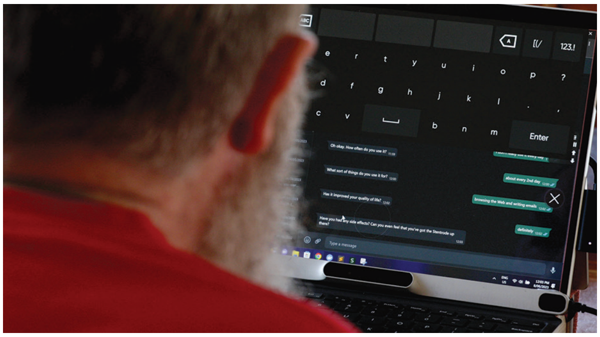

RODNEY GORHAM

Gorham can use his Synchron implant to control not just a computer, but also smart devices in his home like an air conditioner, fan, and smart speaker.

Another way to reduce cognitive load is to detect high-order signals of intent in neural data rather than low-level motor commands, says Florian Solzbacher, cofounder and chief scientific officer of Blackrock Neurotech. For instance, rather than manually navigating to an email app and typing, the user could simply think about sending an email and the system would then open it with content already prepopulated, he says.

Durability may prove a thornier problem to solve, UChicago's Downey says. Current implants last around a decade—well short of a lifelong solution. And with limited real estate in the brain, replacement is possible only once or twice, he says.

Rapid technological progress also raises difficult decisions about whether to get a BCI implant now or wait for a more advanced device. This was a major concern for Gorham's wife, Caroline. "I was hesitant. I didn't want him to go on the trial but maybe a future one," she says. "It was my fear of missing out on future upgrades."

SOME EXECUTIVES HAVE raised the prospect of BCIs eventually becoming consumer devices. Neuralink founder Elon Musk has been particularly vocal, suggesting that the company's implants could replace smartphones, let people save and replay memories, or even achieve "symbiosis" with AI.

This kind of talk inspires mixed feelings in users. The hype brings visibility and funding, says Beggin, but could divert attention from medical users' needs. BCI user Copeland worries that consumer branding could strip the devices of insurance coverage and that rising demand may make it harder to access qualified surgeons.

There are also concerns about how data collected by BCI companies will be handled if the devices go mainstream. As a trial participant, Arbaugh says he's comfortable signing away his data rights to advance the technology, but he thinks stronger legal protections will be needed in the future. "Does that data still belong to Neuralink? Does it belong to each person? And can that data be sold?" he asks.

Blackrock Neurotech's Solzbacher says the company remains focused on the medical applications of the technology. But he also believes it is building a "universal interface to any kind of a computerized system" that may have broader applications in the future. And he says the company owes it to users not to limit them to a bare-bones assistive technology. "Why would somebody who's got a medical condition want to get less than something that somebody who's able-bodied would possibly also take?" says Solzbacher.

The ever-optimistic Imbrie heartily agrees. Medical devices are invariably expensive, he says, but targeting consumer applications could push companies to keep devices simple and affordable while continuing to add features. "I truly believe that making it a consumer-available product will just enhance the product's capabilities for the medical field," he says.

Imbrie is on a mission to refocus the conversation around BCIs on the positives. While concerns about risks are valid, he worries that the alarming language often used to describe brain implants discourages people from volunteering for trials that could help them.

"I remember laying there in the bed and not being able to move," he says, "and it was really dehumanizing having to ask someone to do everything for you. As humans, we want to be independent." ■

RODNEY DECKER

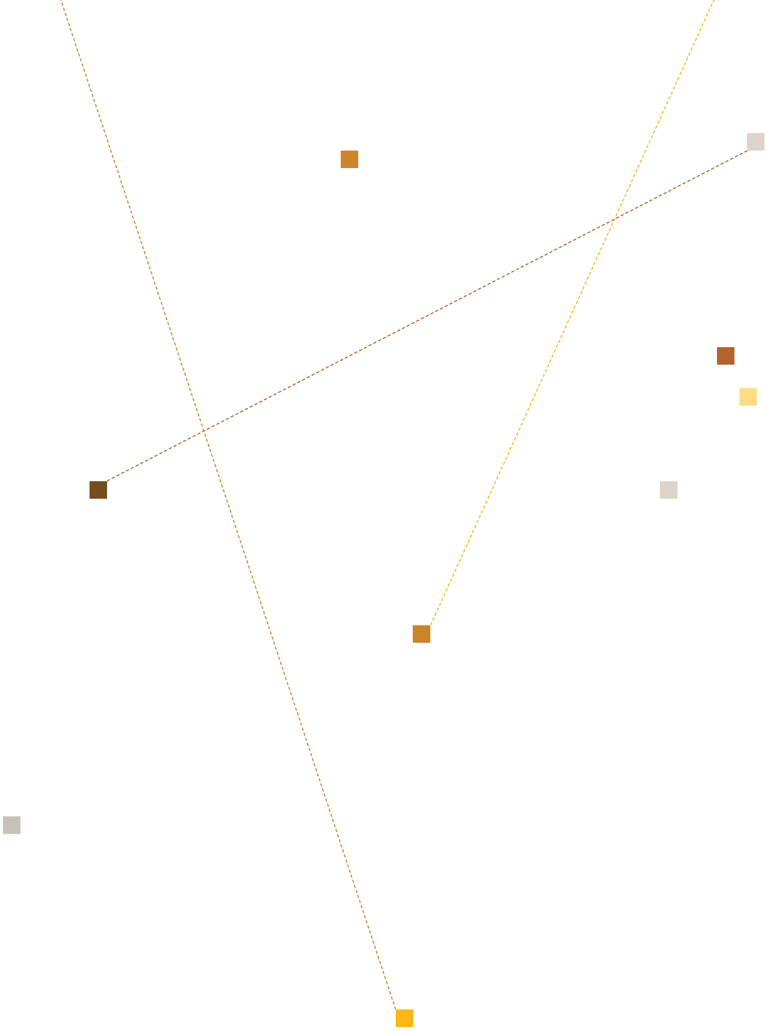

THE STEALTH SIGNALS BYPASSING IRAN’S INTERNET BLACKOUT

Files hidden in satellite TV broadcasts keep information flowing

BY EVAN ALIREZA FIROOZI

Illustration by Mark Harris

You Are Not Connected to the Internet

This page can't be displayed because your computer is currently offline.

On 8 January 2026, the Iranian government imposed a near-total communications shutdown. It was the country's first full information blackout: For weeks, the internet was off across all provinces while services including the government-run intranet, VPNs, text messaging, mobile calls, and even landlines were severely throttled. It was an unprecedented lockdown that left more than 90 million people cut off not only from the world, but from one another.

Since then, connectivity has never fully returned. Following U.S. and Israeli airstrikes in late February, Iran again imposed near-total restrictions, and people inside the country again saw global information flows dry up.

Iran's internet infrastructure is more centralized than in other parts of the world, making it easier for the government to restrict the flow of information.

PREVIOUS PAGE, SOURCE IMAGES: GETTY IMAGES; ISTOCK; THIS PAGE: MORTEZA NIKOUBAZL/NURPHOTO/GETTY IMAGES

The original January shutdown came amid nationwide protests over the deepening economic crisis and political repression, in which millions of people chanted antigovernment slogans in the streets. While Iranian protests have become frequent in recent years, this was one of the most significant uprisings since the Islamic Revolution in 1979. The government responded quickly and brutally. One report put the death toll at more than 7,000 confirmed deaths and more than 11,000 under investigation. Many sources believe the death toll could exceed 30,000.

Thirteen days into the January shutdown, we at NetFreedom Pioneers (NFP) turned to a system we had built for exactly this kind of moment—one that sends files over ordinary satellite TV signals. During the national information vacuum, our technology, called Toosheh, delivered real-time updates into Iran, offering a lifeline to millions starved of trusted information.

I JOINED NetFreedom Pioneers, a nonprofit focused on anticensorship technology, in 2014. Censorship in Iran was a defining feature of my youth in the 1990s. After the Islamic Revolution, most Iranians began to lead double lives—one at home, where they could drink, dance, and choose their clothing, and another in public, where everyone had to comply with stifling government laws.

My first experience with secret communications was when I was five and living in the small city of Fasa in southern Iran. My uncle brought home a satellite dish—dangerously illegal at the time—that allowed us to tune into 12 satellite channels. My favorite was Cartoon Network. Then, during my teenage years, this same uncle introduced me to the internet through dial-up modems. I remember using Yahoo Mail with its 4 megabytes of storage, reading news from around the world, and learning about the Chandra X-ray telescope from NASA's website.

Many Iranian citizens have free-to-air satellite dishes, like the ones on this apartment building in Tehran, and can thus download Toosheh transmissions, giving them a lifeline during internet blackouts.

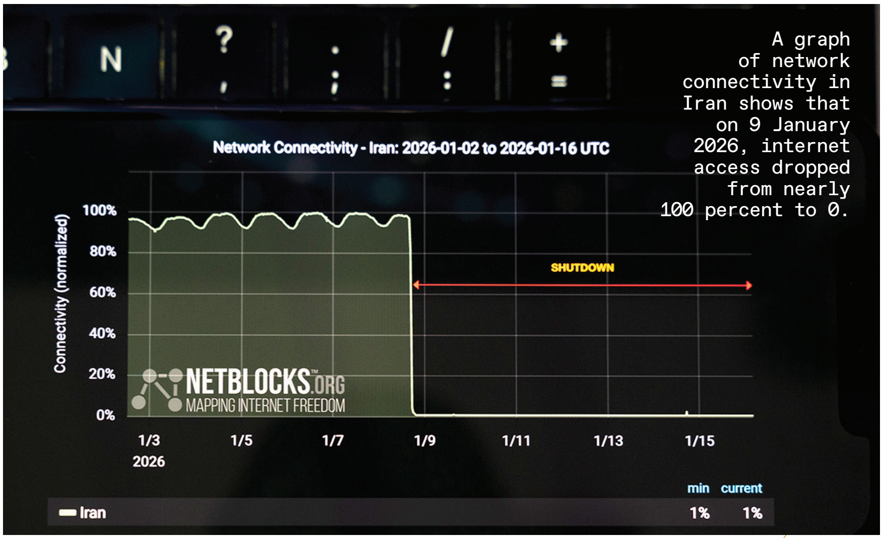

A graph of network connectivity in Iran shows that on 9 January 2026, internet access dropped from nearly 100 percent to 0.

LEFT: MORTEZA NIKOUBAZL/NURPHOTO/GETTY IMAGES;
RIGHT: SAMUEL BOIVIN/NURPHOTO/GETTY IMAGES

That openness didn't last. As internet use spread in the early 2000s, the Iranian government began reshaping the network itself. Unlike the highly distributed networks in the United States or Europe, where thousands of providers exchange traffic across many independent routes, Iran's connection to the global internet is relatively centralized. Most international traffic passes through a small number of gateways controlled by state-linked telecom operators. That architecture gives authorities unusual leverage: By restricting or withdrawing those connections, they can sharply reduce the country's access to the outside world.

Over the past decade, Iran has expanded this control through what it calls the National Information Network, a domestically routed system designed to keep data inside the country whenever possible. Many government services, banking systems, and local platforms are hosted on this internal network. During periods of unrest, access to the global internet can be throttled or cut off while portions of this domestic network continue to function.

The government began its censorship campaign by redirecting or blocking websites. As internet use grew, it adopted more sophisticated approaches. For example, the Telecommunication Company of Iran uses a technique called deep packet inspection to analyze the content of data packets in real time. This method enables it to identify and block specific types of traffic, such as VPN connections, messaging apps, social media platforms, and banned websites.

ACCESS TO INFORMATION IS NOT ONLY ABOUT NEWS OR POLITICS, BUT ABOUT EXPOSURE TO POSSIBILITIES.

TOOSHEH'S communication workaround builds on a history of satellite TV adoption in Middle Eastern and North African countries. By the early 2000s, satellite dishes were common in Iran; today the majority of households in Iran have access to satellite TV despite its official prohibition.

Unlike subscription services such as DirecTV and Dish Network, "free-to-air" satellite TV broadcasts are unencrypted and can be received by anyone with a dish and receiver—no subscription required. Because the signals are open, users can also capture and store the data they carry, rather than simply watching it live. Tech-savvy people learned that they could use a digital video broadcasting (DVB) card—a piece of hardware that connects to a computer and tunes into satellite frequencies—to transform a personal computer into a satellite receiver. This way, they could watch and store media locally as well as download data from dedicated channels.

Toosheh, named for a Persian word that translates to "knapsack," is the brainchild of Mehdi Yahyanejad, an Iranian-American technologist and entrepreneur. Yahyanejad cofounded NetFreedom Pioneers in 2012. He proposed that the satellite-computer connections enabled by a DVB card could be re-created in software, eliminating the need for specialized hardware. He added a simple digital interface to the software to make it easy for anyone to use. The next breakthrough came when the NFP team developed a new transfer protocol that tricks ordinary satellite receivers into downloading data alongside audio and video content. Thus, Toosheh was born.

Satellite TV uses a file system called an MPEG transport stream that allows multiple audio, video, or data layers to be packaged into a single stream file. When you tune in to a satellite channel and select an audio option or closed captions, you're accessing data stored in different parts of this stream. The NFP team's insight was that, by piggybacking on one of these layers, Toosheh could send an MPEG stream that included documents, videos, and more.

A satellite receiver can't tell the difference between our data and normal satellite audio and video data since it only

WE AIM NOT JUST TO BUILD A TOOL FOR CENSORSHIP CIRCUMVENTION, BUT TO REDEFINE ACCESS ITSELF.

"sees" the MPEG streams, not what's encoded on them. This means the data can be downloaded and read, watched, and saved on local devices such as computers, smartphones, or storage devices. What's more, the system is entirely private: No one can detect whether someone has received data through Toosheh; there are no traceable logs of user activity.

Toosheh doesn't provide internet access, but rather delivers curated data through satellite technology. The fundamental distinction lies in the way users interact with the system. Unlike traditional internet services, where you type a request into your browser and receive data in response, Toosheh operates more like a combination of radio and television, presenting information in a magazine-like format. Users don't make requests; instead, they receive 1 to 5 gigabytes of prepackaged, carefully selected data.

During this year's internet blackout, we distributed official statements from Iranian opposition leader Crown Prince Reza Pahlavi and the U.S. government. We provided first-aid tutorials for medics and injured protesters. We sent uncensored news reports from BBC Persian, Iran International, IranWire, VOA Farsi, and others. We also shared critical software packages including anticensorship and antisurveillance tools, along with how-to guides to help people securely connect to Starlink satellite terminals, allowing them to stay protected and anonymous as they sent their own communications.

BECAUSE TOOSHEH relies on one-way satellite broadcasts, it evades the usual tactics governments use to block internet access. However, it remains vulnerable to satellite signal jamming.

The Iranian government is notorious for deploying signal jamming, especially in larger cities. In 2009, the government used uplink interference, which attacks the satellite in orbit by beaming strong noise in the frequency of the satellite's receiver. This makes it impossible for the satellite to distinguish the information it's supposed to receive. However, because this type of attack temporarily disables the entire satellite, Iran was threatened with international sanctions and in 2012 stopped using the method.

The internet shutdown in Iran continued after the attacks by Israel and the United States began in late February, preventing Iranians from communicating with the outside world and with one another.

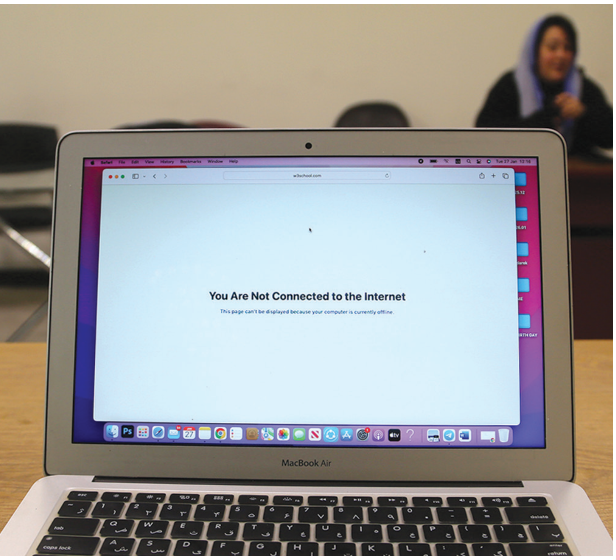
You Are Not Connected to the Internet
This page can't be displayed because your computer is currently offline.
MacBook Air

FATEMEH BAHRAMI/ANADOLU/GETTY IMAGES

The current method, called terrestrial jamming, uses antennas installed at higher elevations than the surrounding buildings to beam strong noise over a specific area in the frequency range of household receivers. This attack is effective in keeping some of the packets from arriving and damaging others, effectively jamming the transmission. But it's short-range and requires significant power, so it's impossible to implement nationwide. There are always people somewhere who can still watch TV, download from Toosheh, or tune into a satellite radio despite the jamming. Even so, we wanted a workaround that would keep our transmissions broadly accessible.

NFP's solution was to add redundancy, similar in principle to a data-storage technique called RAID (redundant array of independent disks). Instead of sending each piece of data once, we send extra information that allows missing or corrupted packets to be reconstructed. Under normal circumstances, we often use 5 percent of our bandwidth for this redundancy. During periods of active jamming, we increase that to as much as 25 to 30 percent, improving the chances that users can recover complete files despite interference.

TOOSHEH INITIALLY came online in 2015 in Iran and Afghanistan. Its full potential, however, was first realized during the 2019 protests in Iran, which saw the most widespread internet shutdown prior to the blackout this year. *Wired* called the 2019 shutdown "the most severe disconnection" tracked by NetBlocks in any country in terms of its "technical complexity and breadth." Our technology helped thousands of people stay informed. We sent crucial local updates, legal-aid guides, digital security tools, and independent news to satellite receivers all over the country, seeing a sixfold increase in our user base.

When that wave of protests subsided, the government allowed some communication services to return. People were again able to access the free internet using VPNs and other antifilter software that allowed them to bypass restrictions. Toosheh then became a public access point for news, educational material, and entertainment beyond government filtering.

Toosheh's impact is often personal. A traveling teacher in western Iran told NFP that he regularly distributed Toosheh files to students in remote villages. One package included footage of female athletes competing in the Olympic Games, something never broadcast in Iran. For one young girl, it was the first time she realized women could compete professionally in sports. That moment underscores a broader truth: Access to information is not only about news or politics, but about exposure to possibilities.

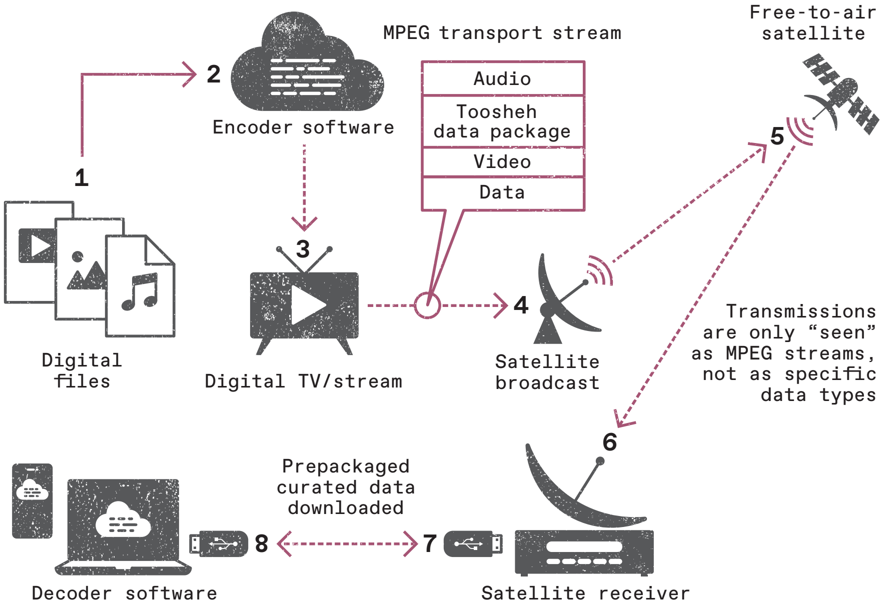

HOW TOOSHEH WORKS: At NetFreedom Pioneers, content curators pull together files—news articles, videos, audio, and software [1]. Toosheh's encoder software [2] compresses the files into a bundle, in .ts format, creating an MPEG transport stream [3]. From there, it's uploaded to a server for transmission [4] via a free-to-air TV channel on a Yahsat satellite that's positioned over the Middle East to provide regional coverage [5]. Satellite receivers [6] directly capture the data streams, which are downloaded to computers, smartphones, and other devices, and decoded by Toosheh software [8].

UNLIKE INTERNET-BASED systems, Toosheh has operational costs that remain constant regardless of the number of users. A single TV satellite in geostationary earth orbit, deployed and maintained by an international company such as Eutelsat, can broadcast to an entire continent with no increase in cost to audiences. What's more, the startup cost for users isn't high: A satellite dish and receiver in Iran costs less than US $50, which is affordable to many. And it costs nothing for people to use Toosheh's service and receive its files.

However, operating the service is costly: NetFreedom Pioneers pays tens of thousands of dollars a month for satellite bandwidth. We had received funding from the U.S. State Department, but in August of 2025, that funding ended, forcing us to suspend services in Iran.

Then the December protests happened, and broadcasting to Iran became an urgent priority. To turn Toosheh back on, we needed roughly $50,000 a month. With the support of a handful of private donors, we were able to meet these costs and sustain operations in Iran for a few months, though our future there and elsewhere is uncertain.

TOOSHEH'S REVIVAL in Iran came alongside our ongoing support for deployments of Starlink, a satellite internet service that allows users to connect directly to satellites rather than relying on domestic networks, which the government can shut down. Unlike Toosheh's one-way broadcasts, Starlink provides full two-way internet access, enabling users to send messages, upload videos, and communicate with the outside world.

In 2022, we started gathering donations to buy Starlink terminals for Iran. We have delivered more than 300 of the roughly 50,000 there, enabling citizens to send encrypted updates and videos to us from inside the country. Because the technology is banned by the government, access remains limited and carries risk; Iranian authorities have recently arrested Starlink users and sellers. And unlike Toosheh's receive-only broadcasts, Starlink terminals transmit signals back to orbit, creating a radio footprint that can potentially be detected.

Looking ahead, we envision Toosheh becoming a foundational part of global digital resilience. It is uncensored, untraceable, and resistant to government shutdowns. Because Toosheh is downlink only, it can sometimes feel hard to explain the value of this technology to those living in the free world, those accustomed to open internet access. Yet, people living under censorship have few other choices when there's a digital blackout.

Currently, NFP is developing new features like intelligent content curation and automatically prioritizing data packages based on geographic or situational needs. And we're experimenting with local sharing tools that allow users who receive Toosheh broadcasts to redistribute those files via Wi-Fi hotspots or other offline networks, which could extend the system's reach to disaster zones, conflict areas, and climate-impacted regions where infrastructure may be destroyed.

We're also looking at other use cases. Following the Taliban's return to power in Afghanistan, NetFreedom Pioneers designed a satellite-based system to deliver educational materials. Our goal is to enable private, large-scale distribution of coursework to anyone—including the girls who are banned from Afghanistan's schools. The system is technically ready but has yet to secure funding for deployment.

We aim not just to build a tool for censorship circumvention, but to redefine access itself. Whether in an Iranian city under surveillance, a Guatemalan village without internet, or a refugee camp in East Africa, Toosheh offers a powerful and practical model for delivering vital information without relying on vulnerable or expensive networks.

Toosheh is a reminder that innovation doesn't have to mean complexity. Sometimes, the most transformative ideas are the simplest, like delivering data through the sky, quietly and affordably, into the hands of those who need it most. ■

Illustration by Chris Philpot

Smart cities will treat interference as a new kind of calculator

BY ANA I. PÉREZ-NEIRA & ALPHAN SAHIN

Teaching Radio Waves to Compute

Illustration by Chad Hagen

Picture a highway with networked autonomous cars driving along it. On a serene, cloudless day, these cars need only exchange thimblefuls of data with one another. Now picture the same stretch in a sudden snow squall: The cars rapidly need to share vast amounts of essential new data about slippery roads, emergency braking, and changing conditions.

These two very different scenarios involve vehicle networks with very different computational loads. Eavesdropping on network traffic using a ham radio, you wouldn't hear much static on the line on a clear, calm day. On the other hand, sudden whiteout conditions on a wintry day would sound like a cacophony of sensor readings and network chatter.

Normally this cacophony would mean two simultaneous problems: congested communications and a rising demand for computing power to handle all the data. But what if the network itself could expand its processing capabilities with every rising decibel of chatter and with every sensor's chirp?

Traditional wireless networks treat communication as separate from computation. First you move data, then you process it. However, an emerging new paradigm called over-the-air computation (OAC) could fundamentally change the game. First proposed in 2005 and recently developed and prototyped by a number of teams around the world, including ours, OAC combines communication and computation into a single framework. This means that an OAC sensor network—whether shared among autonomous vehicles, Internet-of-Things sensors, smart-home devices, or smart-city infrastructure—can carry some of the network's computing burden as conditions demand.

The idea takes advantage of a basic physical fact of electromagnetic radiation: When multiple devices transmit simultaneously, their wireless signals naturally combine in the air. Normally, such cross talk is seen as interference, which radios are designed to suppress—especially digital radios with their error-correcting schemes and inherent resistance to low-level noise.

But if we carefully design the transmissions, cross talk can enable a wireless network to directly perform some calculations, such as a sum or an average. Some prototypes today do this with analog-style signaling on otherwise digital radios, so that the superimposed waveforms represent numbers that have been added before digital signal processing takes place.

Researchers are also beginning to explore digital, over-the-air computation schemes, which embed the same ideas into digital formats, ultimately allowing the prototype schemes to coexist with today's digital radio protocols. These various over-the-air computation techniques can help networks scale gracefully, enabling new classes of real-time, data-intensive services while making more efficient use of wireless spectrum.

OAC, in other words, turns signal interference from a problem into a feature, one that can help wireless systems support massive growth.

FOR DECADES, ENGINEERS DESIGNED radio communications protocols with one overriding goal: to isolate each signal and recover each message cleanly. Today's networks face a different set of pressures. They must coordinate large groups of devices on shared tasks—such as AI model training or combining disparate sensor readings, also known as sensor fusion—while exchanging as little raw data as possible, to improve both efficiency and privacy. For these reasons, a new approach to transmitting and receiving data may be worth considering, one that doesn't rely on collecting and storing every individual device's contributions.

By turning interference into computation, OAC transforms the wireless medium from a contested battlefield into a collaborative workspace. This paradigm shift has far-reaching consequences: Signals no longer compete for isolation; they cooperate to achieve shared outcomes. OAC cuts through layers of digital processing, reduces latency, and lowers energy consumption.

Even very simple operations, such as addition, can be the building blocks of surprisingly powerful computations. Many complex processes can be broken down into combinations of simpler pieces, much like how a rich sound can be re-created by combining a few basic tones. By carefully shaping what devices transmit and how the result is interpreted at the receiver, the wireless channel running OAC can carry out other calculations beyond addition. In practice, this means that with the right design, wireless signals can compute a number of key functions that modern algorithms rely on.

For instance, many key tasks in modern networks don't require the logging and storage of every individual network transmission. Rather, the goal is instead to infer properties about aggregate patterns of network traffic—reaching agreement or identifying what matters most about the traffic. Consensus algorithms rely on majority voting to ensure reliable decisions, even when some devices fail. Artificial intelligence systems depend on matrix reduction and simplification operations such as "max pooling" (keeping only peak values) to extract the most useful signals from noisy data.

In smart cities and smart grids, what matters most is often not individual readings but distribution. How many devices report each traffic condition? What is the range of demand across neighborhoods? These

The traditional approach to the problem Consider five connected vehicles traveling within sight of one another. Each car reports its speed to the network. In this example, the speeds are slow, medium, and fast. Using existing standards, all five connected cars must independently track and count all incoming signals. Even in this very simplified case, the network is already congested.

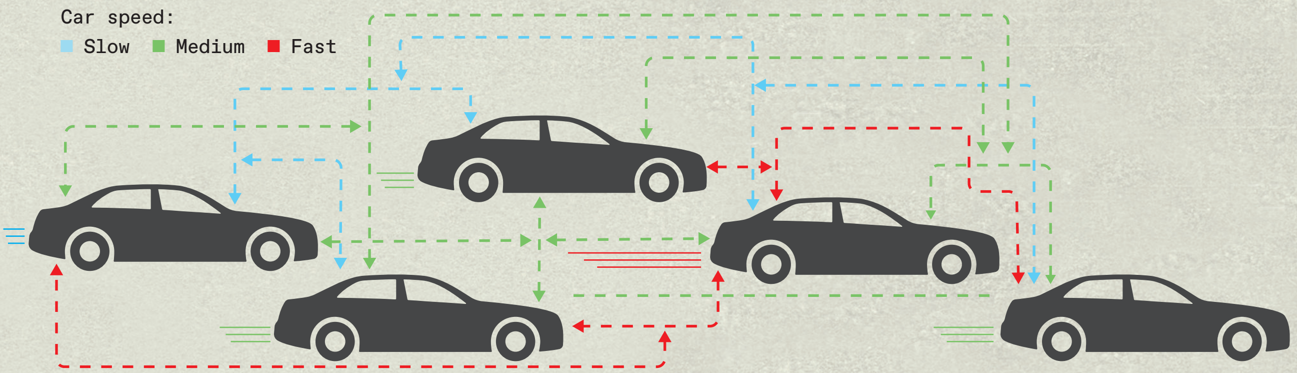

The over-the-air computation (OAC) approach Using over-the-air computation, all five cars transmit their speeds simultaneously. Vehicles reporting the same speed share the same channel; their signals merely combine over the air.

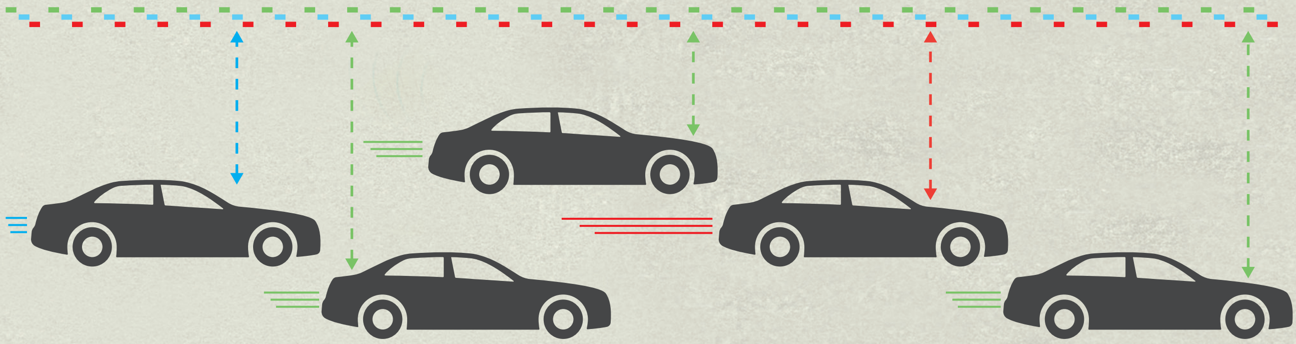

The OAC result One transmission yields the full picture: One car is going slow; three are traveling at medium speed; and one vehicle is moving fast. The majority condition is immediately identified—with no individual vehicle data shared or processed.

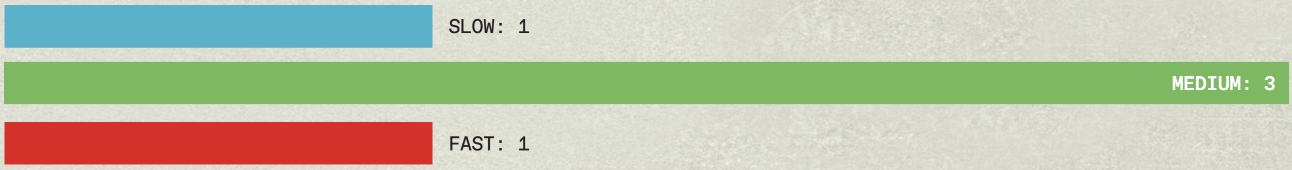

are histogram questions—summaries of the device counts per category.

With type-based multiple access (TBMA), an over-the-air computation method we use, devices reporting a given condition transmit together over a shared channel. Their signals add up, and the receiver sees only the total signal strength per category. In a single transmission, the entire histogram emerges without ever identifying individual devices. And the more devices there are, the better the estimate. The result is greater spectrum efficiency, with lower latency and scalable, privacy-friendly operations—all from letting the wireless medium do the aggregating and counting.

It's easy to imagine how analog values transmitted over the air could be summed via superposition.

Illustration by Mark Montgomery

The amplitudes from different signals add together, so the values those amplitudes represent also simply add together. The more challenging question concerns preserving that additive magic, but with digital signals.

Here's how OAC does it. Consider, for instance, one TBMA approach for a network of sensors that gives each possible sensor reading its own dedicated frequency channel. Every sensor on the network that reads "4" transmits on frequency four; every sensor that reads "7" transmits on frequency seven. When multiple devices share the same reading, their amplitudes combine. The stronger the combined signal at a given frequency, the more devices there are reporting that particular value.

A receiver equipped with a bank of filters tuned to each frequency reads out a count of votes for every possible sensor value. In a single simultaneous transmission, the whole network has reported its state.

95%

Accuracy achieved in training an AI algorithm to recognize images using over-the-air computation

It might seem paradoxical—digital computation riding atop what appears to be an analog physical effect. But this is also true of all "digital" radio. A Wi-Fi transmitter does not launch ones and zeroes into the air; it modulates electromagnetic waves whose amplitudes and phases encode digital data. The "digital" label ultimately refers to the information layer, not the physics. What makes OAC digital, in the same sense, is that the values being computed—each sensor reading, each frequency-bin count—are discrete and quantized from the start. And because they are discrete, the same error-correction machinery that has made digital communications robust for decades can be applied here too.

Synchronization is where OAC's demands diverge most sharply from digital wireless conventions. Many OAC variants today require something akin to a shared clock at nanosecond precision: Every signal's phase must be synchronized, or the superposition runs the risk of collapsing into destructive interference. While TBMA relaxes this burden a bit—devices need only share a time window—real engineering challenges lie ahead regardless, before over-the-air computation is ready for the mobile world.

OVER-THE-AIR COMPUTATION has in recent years moved from theory to initial proofs-of-concept and network test runs. Our research teams in South Carolina and Spain have built working prototypes that deliver repeatable results—with no cables and no external timing sources such as GPS-locked references. All synchronization is handled within the radios themselves.

Our team at the University of South Carolina (led by Sahin) started with off-the-shelf software-defined radios—Analog Devices' Adalm-Pluto. We modified the devices' field-programmable gate array hardware inside each radio so it can respond to a trigger signal transmitted from another radio. This simple hack enabled simultaneous transmission, a core requirement for OAC. Our setup used five radios acting as edge devices and one acting as a base station. The task involved training a neural network to perform image recognition over the air. Our system, whose results we first reported in 2022, achieved a 95 percent accuracy in image recognition without ever moving raw data across the network.

We also demonstrated our initial OAC setup at a March 2025 IEEE 802.11 working group meeting, where an IEEE committee was studying AI and machine learning capabilities for future Wi-Fi standards. As we showed, OAC's road ahead doesn't necessarily require reinventing wireless technology. Rather, it can also build on and repurpose existing protocols already in Wi-Fi and 5G.

However, before OAC can become a routine feature of commercial wireless systems, networks must provide finer-tuned coordination of timing and signal power levels. Mobility is a difficult problem, too. When mobile devices move around, phase synchronization degrades quickly, and computational accuracy can suffer. Present-day OAC tests work in controlled lab environments. But making them robust in dynamic, real-world settings—vehicles on highways, sensors scattered across cities—remains a new frontier for this emerging technology.

Both of our teams are now scaling up our prototypes and demonstrations. We are together aiming to understand how over-the-air computation performs as the number of devices increases beyond lab-bench scales. Turning prototypes and test-beds into production systems for autonomous vehicles and smart cities will require anticipating tomorrow's mobility and synchronization problems—and no doubt a range of other challenges down the road.

NODE 14
ADALM

Researchers at the University of South Carolina are developing prototype sensor nodes [right] to begin testing small, room-sized networks [far right] that perform over-the-air computation.

TO REALIZE THE TECHNOLOGICAL ambitions of over-the-air computation, nanosecond timing and exquisite RF signal design will be crucial. Fortunately, recent engineering advances have made substantial progress in both of these fields.

Because OAC demands waveform superposition, it benefits from tight coordination in time, frequency, phase, and amplitude among RF transmitters. Such requirements build naturally on decades of work in wireless communication systems designed for shared access. Modern networks already synchronize large numbers of devices using high-precision timing and uplink coordination.

OAC uses the same synchronization techniques already in cellular and Wi-Fi systems. But to actually run over-the-air computations, more precision still will be needed. Power control, gain adjustment, and timing calibration are standard tools today. We expect that engineers will further refine these existing methods to begin to meet OAC's more stringent accuracy demands.

In some cases, in fact, imperfect timing standards may be all that's needed. Designs and emerging standards in 5G and 6G wireless systems today use clever encoding that tolerates imperfect synchronization. Minor timing errors, frequency drift, and signal overlap can in some cases still work capably within an OAC protocol, we anticipate. Instead of fighting messiness, over-the-air computation may sometimes simply be able to roll with it.

Another challenge ahead concerns shifting processing to the transmitter. Instead of the receiver trying to clean up overlapping signals, a better and more efficient approach would involve each transmitter fixing its own signal before sending. Such "pre-compensation" techniques are already used in MIMO technology (multiantenna systems in modern Wi-Fi and cellular networks). OAC would just be repurposing techniques that have already been developed for 5G and 6G technologies.

OAC will evolve from lab prototype to standardized wireless capability through the 2020s and into the decade ahead.

ALPHAN SAHIN/UNIVERSITY OF SOUTH CAROLINA

Materials science can also help OAC efforts ahead. New generations of reconfigurable intelligent surfaces shape signals via tiny adjustable elements in the antenna. The surfaces catch radio signals and reshape them as they bounce around. Reconfigurable surfaces can strengthen useful signals, eliminate interference, and synchronize wavefront arrivals that would otherwise be out of sync. OAC stands to benefit from these and other emerging capabilities that intelligent surfaces will provide.

At the system level, OAC will represent a fundamental shift in wireless network system design. Wireless engineers have traditionally tried to avoid designing devices that transmit at the same time. But over-the-air systems will flip the old, familiar design standards on their head.

One might object that OAC stands to upend decades of existing wireless signal standards that have always presumed data pipes to be data pipes only—not microcomputers as well. Yet we do not anticipate much difficulty merging OAC with existing wireless standards. In a sense, in fact, the IEEE 802.11 and 3GPP (3rd Generation Partnership Project) standards bodies have already shown the way.

A network can set aside certain brief time windows or narrow slices of bandwidth for over-the-air computation, and use the rest for ordinary data. From the radio's point of view, OAC just becomes another operating mode that is turned on when needed and left off the rest of the time.

Over the past decade, both the IEEE and 3GPP have integrated once-experimental technologies into their wireless standards—for example, millimeter-wave mobile communications, multiuser MIMO, beamforming, and network slicing—by defining each new technological advance as an optional feature. OAC, we suggest, can also operate alongside conventional wireless data traffic as an optional service. Because OAC places high demands on timing and accuracy, networks will need the ability to enable or disable over-the-air computation on a per-application basis.

With continued progress, OAC will evolve from lab prototype to standardized wireless capability through the 2020s and into the decade ahead. In the process, the wireless medium will transform from a passive data carrier into an active computational partner—providing essential infrastructure for the real-time intelligent systems that future wireless technologies will demand.

So on that snowy highway sometime in the 2030s, vehicles and sensors won't wait for permission to think together. Using the emerging over-the-air computation protocols that we're helping to pioneer, simultaneous computation will be the new default. The networks will work as one. ■

IEEE *Access*®

Multidisciplinary ⋮ Rapid Review ⋮ Open Access Journal

Harness the publishing power of IEEE *Access*.®

- **Receive high-quality, rigorous peer review** in only 4 to 6 weeks
- **Reach millions of global users** through the IEEE *Xplore*® digital library by publishing open access
- **Submit multidisciplinary articles** that may not fit in narrowly focused journals

Learn more at
ieeeaccess.ieee.org

IEEE

香港科技大学（广州）
THE HONG KONG
UNIVERSITY OF SCIENCE AND
TECHNOLOGY (GUANGZHOU)

Dean of Systems Hub

The Hong Kong University of Science and Technology (Guangzhou), located in Guangzhou, China, is seeking a visionary leader as Dean of Systems Hub, which consists of four Thrust Areas: **Bioscience & Biomedical Engineering, Intelligent Transportation, Robotics & Autonomous Systems, and Smart Manufacturing.**

Responsibilities and Qualification

The Dean holds both an executive and faculty appointment, serves as a key member of the University's senior administrative team and is expected to provide dynamic strategic leadership to advance the Hub's academic mission, foster a collegial and inclusive culture, and drive excellence in interdisciplinary education, research, and innovation.

- **Academic Credentials:** Be distinguished scholars with international recognition for their academic excellence in research and teaching in one of the Hub's disciplines.
- **Leadership Experience:** A proven record of successful leadership and effective administrative management.
- **Communication Skills:** Able to engage and build relationships with various stakeholders, including faculty, students from varied cultural backgrounds, and the wider community.
- **Strategic Thinking:** Able to anticipate emerging trends and best practices in science, engineering and to translate these insights into actionable plans.

Remuneration and Application Procedure

This position is a full-time appointment, typically for an initial term of three years, renewable by mutual agreement for another three-year term. Competitive remuneration at international standards, along with generous fringe benefits including medical and housing support, will be provided.

For more information and to apply, please visit:
https://career.hkust-gz.edu.cn/en/job-detail?jobId=1242

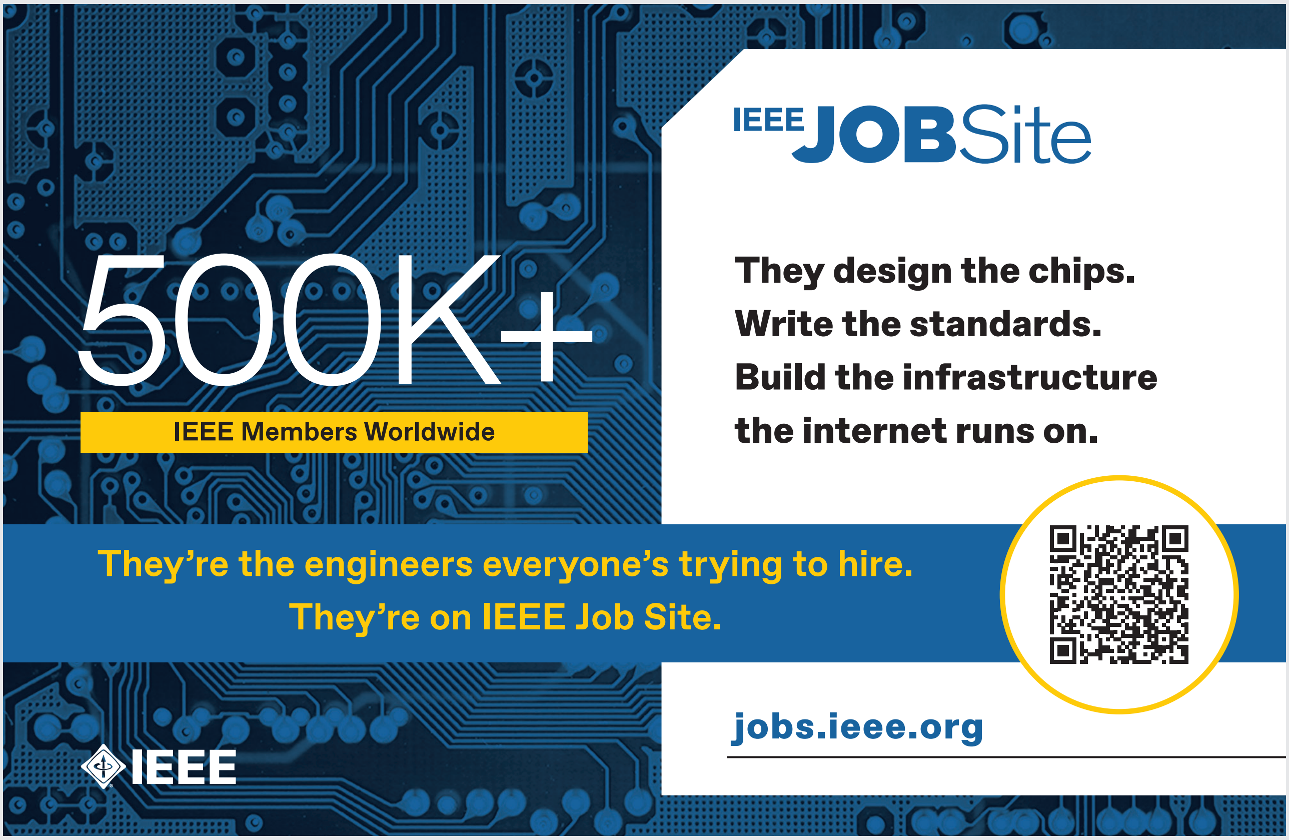
IEEE JOBSite

500K+
IEEE Members Worldwide

They design the chips.
Write the standards.
Build the infrastructure the internet runs on.

They're the engineers everyone's trying to hire.
They're on IEEE Job Site.

jobs.ieee.org

IEEE

BY ALLISON MARSH

Past Forward

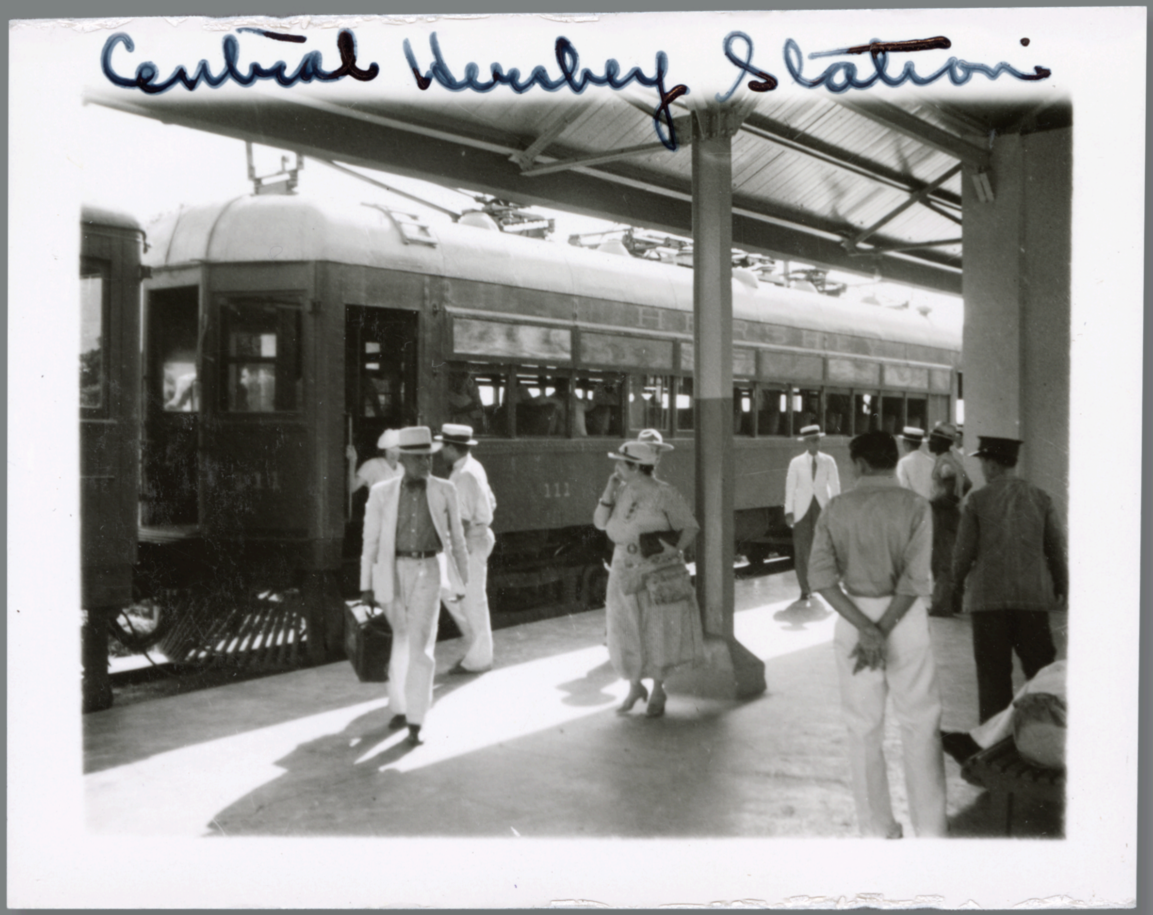

In this 1936 photo, passengers wait at the Central Hershey station, one of the main stops on the Hershey Electric Railway in Cuba.

HERSHEY COMMUNITY ARCHIVES

This Chocolate Empire Ran on Electric Rails

The U.S. chocolate magnate Milton S. Hershey understood the value of vertical integration. So when a global sugar shortage prompted Hershey to begin purchasing sugarcane fields in Cuba in 1916, it naturally made sense to also invest in a Cuban sugar mill, refinery, power plant, and railroad to transport the processed sugar to port. Many of Hershey's Cuban employees lived in well-maintained company towns.

At first, steam locomotives burning oil or coal ran on a single track between Havana and the northern port of Matanzas. But imported fuel was pricey, and so the company began electrifying the line. By 1920, a fleet of seven four-motor, 1,200-volt DC locomotives were hauling carloads of sugarcane and general freight, with additional trains providing passenger service.

Hershey sold its Cuban operation in 1946, and the sugar refinery closed in 2002. Passenger service on the Hershey Electric Railway continued sporadically, despite frequent power outages and damaged infrastructure, and catching a glimpse of—and perhaps a ride on—the line became an aspiration for die-hard railfans. ■

FOR MORE ON HERSHEY'S CUBAN ELECTRIC RAILROAD, see spectrum.ieee.org/pastforward-may2026

IEEE Member
Discounts Program

Unlock Exclusive Savings as an IEEE member.

Get more value from your membership.

Along with staying current on the latest developments, connecting with peers, and accessing high-quality resources, IEEE Membership also gives you exclusive offers and discounts on a wide range of products and services.

Visit **ieee.org/discounts** to explore member savings on brand discounts and country-specific insurance plan options within the IEEE Member Discounts Program.

92-FA-0009 03/26

IEEE
Advancing Technology for Humanity

Tomorrow’s planes. Engineered today with MATLAB and Simulink.

JetZero’s departure from conventional fuselage and tail design is purposeful: reduce drag, increase lift, and deliver meaningful fuel reductions.

mathworks.com/engineered-systems

MathWorks®
Accelerating the pace of engineering and science

© The MathWorks